

US007395405B2

(12) **United States Patent**
Anderson et al.

(10) **Patent No.:** **US 7,395,405 B2**
(45) **Date of Patent:** **Jul. 1, 2008**

(54) **METHOD AND APPARATUS FOR SUPPORTING ADDRESS TRANSLATION IN A VIRTUAL MACHINE ENVIRONMENT**

4,319,323 A 3/1982 Ermolovich et al.
4,347,565 A 8/1982 Kaneda et al.
4,366,537 A 12/1982 Heller et al.

(75) Inventors: **Andrew V. Anderson**, Hillsboro, OR (US); **Alain Kägi**, Portland, OR (US)

(Continued)

(73) Assignee: **Intel Corporation**, Santa Clara, CA (US)

FOREIGN PATENT DOCUMENTS

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 447 days.

DE 4217444 12/1992

(Continued)

(21) Appl. No.: **11/045,524**

OTHER PUBLICATIONS

(22) Filed: **Jan. 28, 2005**

Berg, Cliff, "How Do I Create a Signed Applet?", *Dr. Dobb's Journal*, (Aug. 1997), 1-9.

(65) **Prior Publication Data**

US 2006/0174053 A1 Aug. 3, 2006

(Continued)

(51) **Int. Cl.**
G06F 12/00 (2006.01)
G06F 21/00 (2006.01)
G06F 9/26 (2006.01)
G06F 9/34 (2006.01)
G06F 9/455 (2006.01)

Primary Examiner—Donald Sparks
Assistant Examiner—Ryan Bertram
(74) *Attorney, Agent, or Firm*—Blakely, Sokoloff, Taylor & Zafman LLP

(52) **U.S. Cl.** **711/206**; 711/6; 711/202; 711/203; 711/205; 711/207; 711/210; 718/1

(57) **ABSTRACT**

(58) **Field of Classification Search** 711/205
See application file for complete search history.

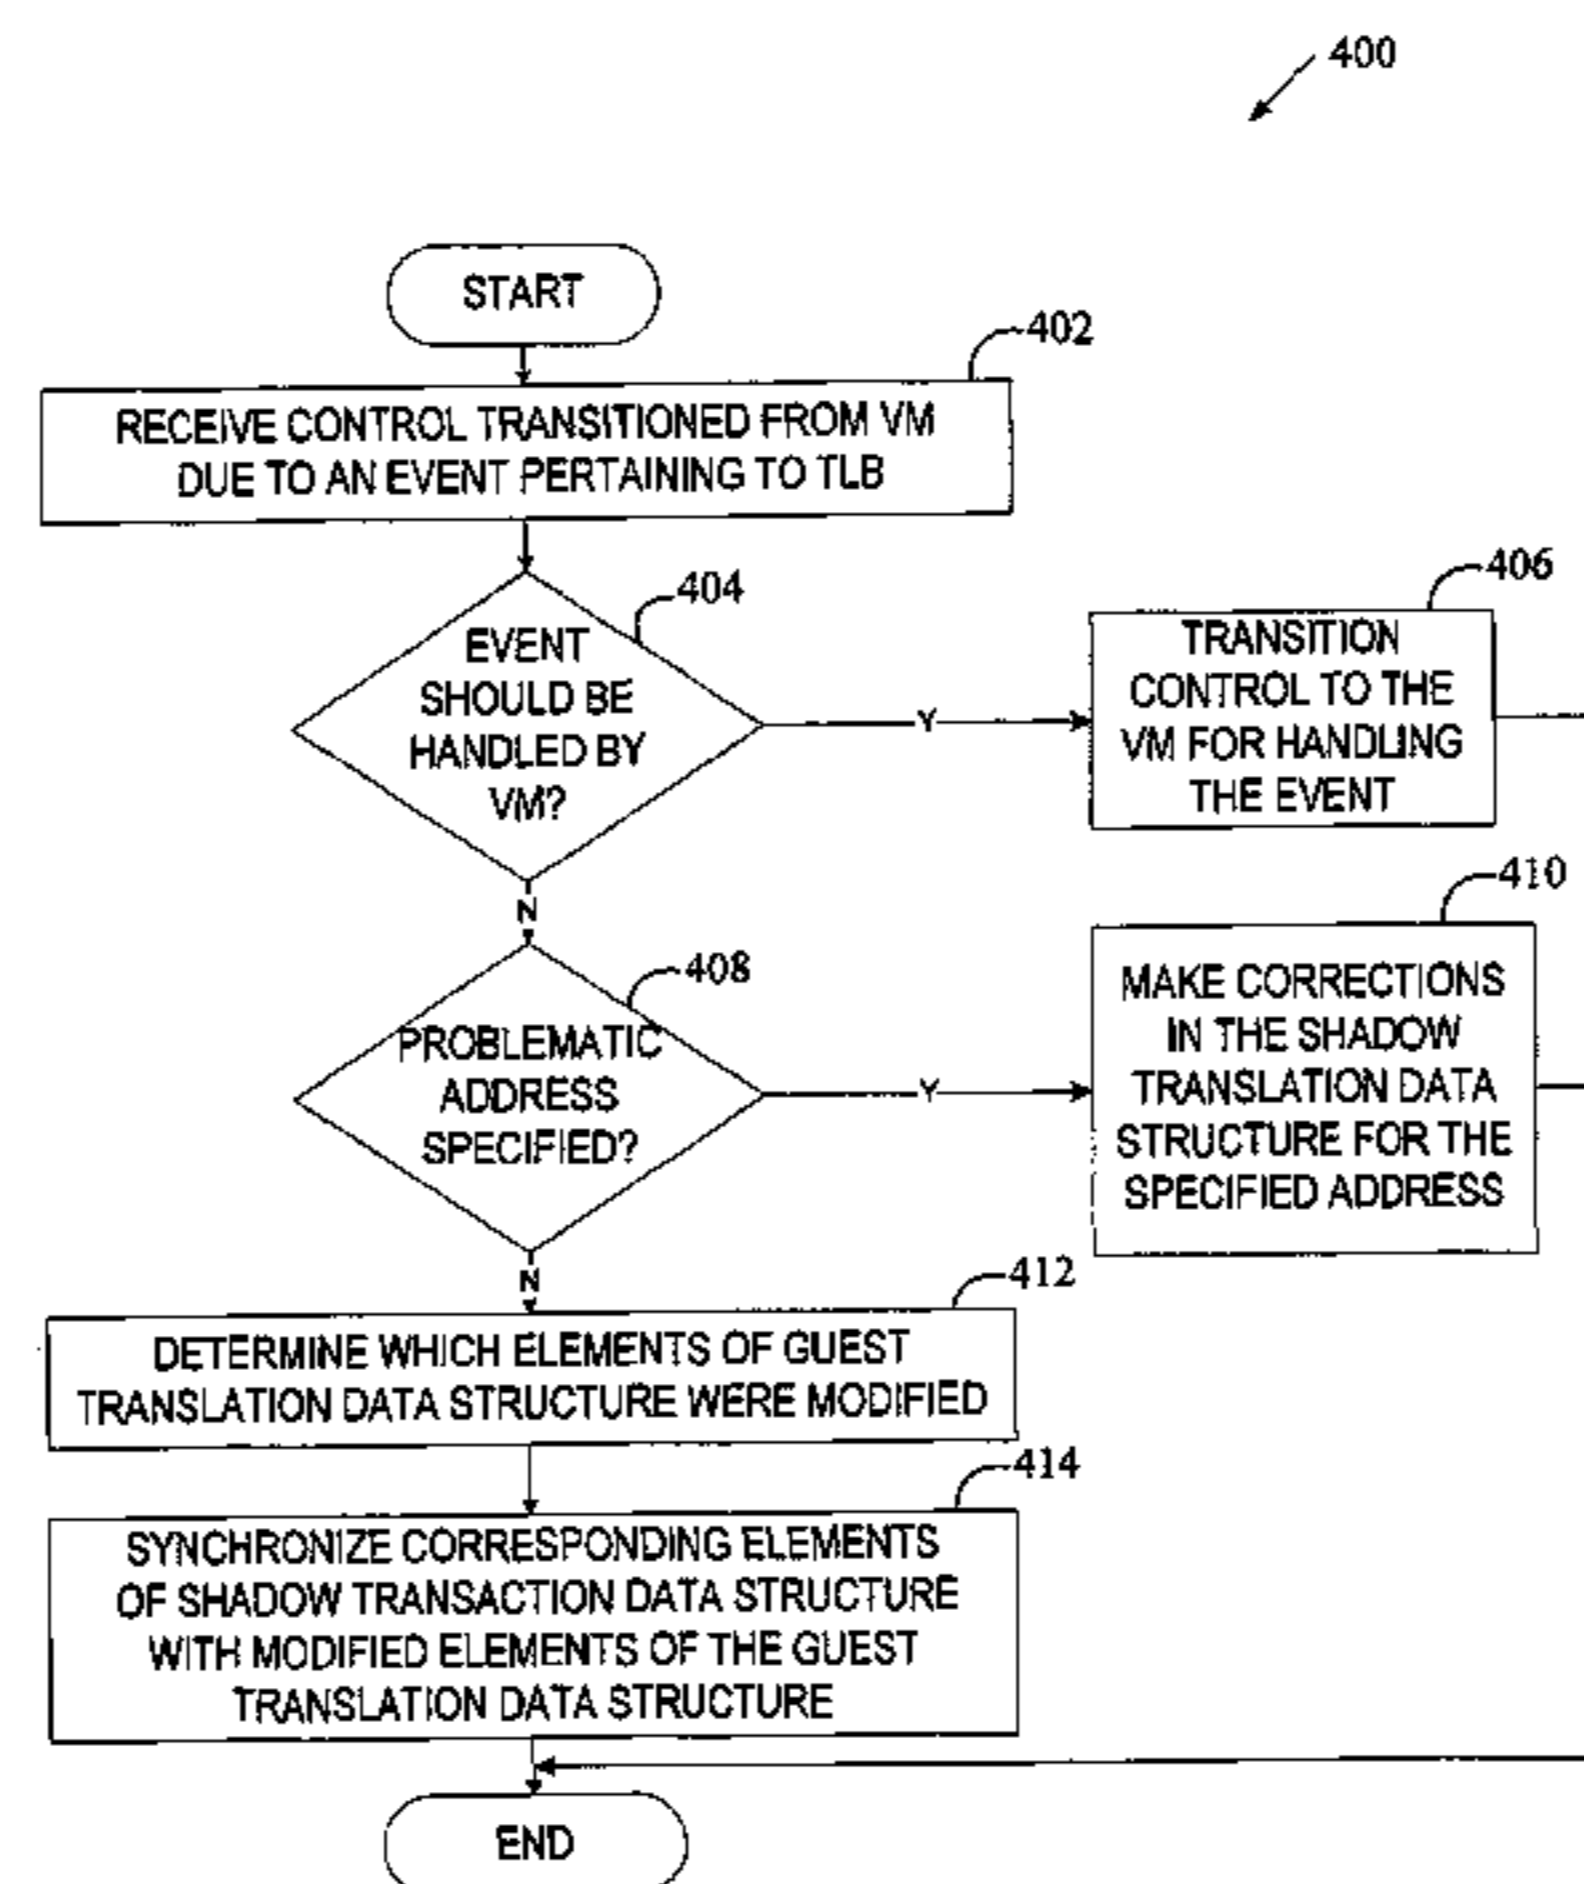
In one embodiment, a method includes receiving control transitioned from a virtual machine (VM) due to a privileged event pertaining to a translation-lookaside buffer (TLB), and determining which entries in a guest translation data structure were modified by the VM. The determination is made based on metadata extracted from a shadow translation data structure maintained by a virtual machine monitor (VMM) and attributes associated with entries in the shadow translation data structure. The method further includes synchronizing entries in the shadow translation data structure that correspond to the modified entries in the guest translation data structure with the modified entries in the guest translation data structure.

(56) **References Cited**

U.S. PATENT DOCUMENTS

3,699,532 A 10/1972 Schaffer et al.
3,996,449 A 12/1976 Attanasio et al.
4,037,214 A 7/1977 Birney et al.
4,162,536 A 7/1979 Morley
4,207,609 A 6/1980 Luiz et al.
4,247,905 A 1/1981 Yoshida et al.
4,276,594 A 6/1981 Morley
4,278,837 A 7/1981 Best
4,307,447 A 12/1981 Provanzano et al.
4,319,233 A 3/1982 Matsuoka et al.

23 Claims, 15 Drawing Sheets



US 7,395,405 B2

Page 2

U.S. PATENT DOCUMENTS					
			5,764,969	A	6/1998 Kahle
			5,796,835	A	8/1998 Saada
			5,796,845	A	8/1998 Serikawa et al.
			5,805,712	A	9/1998 Davis
			5,809,546	A	9/1998 Greenstein et al.
			5,825,875	A	10/1998 Ugon
			5,825,880	A	10/1998 Sudia et al.
			5,835,594	A	11/1998 Albrecht et al.
			5,844,986	A	12/1998 Davis
			5,852,717	A	12/1998 Bhide et al.
			5,854,913	A	12/1998 Goetz et al.
			5,867,577	A	2/1999 Patarin
			5,872,994	A	2/1999 Akiyama et al.
			5,890,189	A	3/1999 Nozue et al.
			5,900,606	A	5/1999 Rigal
			5,901,225	A	5/1999 Ireton et al.
			5,903,752	A	5/1999 Dingwall et al.
			5,919,257	A	7/1999 Trostle
			5,935,242	A	8/1999 Madany et al.
			5,935,247	A	8/1999 Pai et al.
			5,937,063	A	8/1999 Davis
			5,944,821	A	8/1999 Angelo
			5,953,502	A	9/1999 Helbig, Sr.
			5,956,408	A	9/1999 Arnold
			5,970,147	A	10/1999 Davis et al.
			5,978,475	A	11/1999 Schneier et al.
			5,978,481	A	11/1999 Ganesan et al.
			5,987,557	A	11/1999 Ebrahim
			6,014,745	A	1/2000 Ashe
			6,035,374	A	3/2000 Panwar et al.
			6,044,478	A	3/2000 Green
			6,055,637	A	4/2000 Hudson et al.
			6,058,478	A	5/2000 Davis
			6,061,794	A	5/2000 Angelo
			6,075,938	A	6/2000 Bugnion et al.
			6,085,296	A	7/2000 Karkhanis et al.
			6,088,262	A	7/2000 Nasu
			6,092,095	A	7/2000 Maytal
			6,093,213	A	7/2000 Favor et al.
			6,101,584	A	8/2000 Satou et al.
			6,108,644	A	8/2000 Goldschlag et al.
			6,115,816	A	9/2000 Davis
			6,125,430	A	9/2000 Noel et al.
			6,131,166	A	10/2000 Wong-Isley
			6,148,379	A	11/2000 Schimmel
			6,158,546	A	12/2000 Hanson et al.
			6,173,417	B1	1/2001 Merrill
			6,175,924	B1	1/2001 Arnold
			6,175,925	B1	1/2001 Nardone et al.
			6,178,509	B1	1/2001 Nardone
			6,182,089	B1	1/2001 Ganapathy et al.
			6,188,257	B1	2/2001 Buer
			6,192,455	B1	2/2001 Bogin et al.
			6,199,152	B1	3/2001 Kelly et al.
			6,205,550	B1	3/2001 Nardone et al.
			6,212,635	B1	4/2001 Reardon
			6,222,923	B1	4/2001 Schwenk
			6,249,872	B1	6/2001 Wildgrube et al.
			6,252,650	B1	6/2001 Nakaumra
			6,269,392	B1	7/2001 Cotichini et al.
			6,272,533	B1	8/2001 Browne et al.
			6,272,637	B1	8/2001 Little et al.
			6,275,933	B1	8/2001 Fine et al.
			6,282,650	B1	8/2001 Davis
			6,282,651	B1	8/2001 Ashe
			6,282,657	B1	8/2001 Kaplan et al.
			6,292,874	B1	9/2001 Barnett
			6,301,646	B1	10/2001 Hostetter
			6,308,270	B1	10/2001 Guthery et al.
			6,314,409	B2	11/2001 Schneck et al.
			6,321,314	B1	11/2001 Van Dyke
			6,327,652	B1	12/2001 England et al.
			6,330,670	B1	12/2001 England et al.
4,403,283	A	9/1983 Myntti et al.			
4,419,724	A	12/1983 Branigin et al.			
4,430,709	A	2/1984 Schleupen et al.			
4,521,852	A	6/1985 Gutttag			
4,571,672	A	2/1986 Hatada et al.			
4,621,318	A	11/1986 Maeda			
4,759,064	A	7/1988 Chaum			
4,795,893	A	1/1989 Ugon			
4,802,084	A	1/1989 Ikegaya et al.			
4,825,052	A	4/1989 Chemin et al.			
4,907,270	A	3/1990 Hazard			
4,907,272	A	3/1990 Hazard			
4,910,774	A	3/1990 Barakat			
4,975,836	A	12/1990 Hiroswawa et al.			
5,007,082	A	4/1991 Cummins			
5,022,077	A	6/1991 Bealkowski et al.			
5,075,842	A	12/1991 Lai			
5,079,737	A	1/1992 Hackbarth			
5,187,802	A	2/1993 Inoue et al.			
5,230,069	A	7/1993 Brelsford et al.			
5,237,616	A	8/1993 Abraham et al.			
5,255,379	A	10/1993 Melo			
5,287,363	A	2/1994 Wolf et al.			
5,293,424	A	3/1994 Holtey et al.			
5,295,251	A	3/1994 Wakui et al.			
5,317,705	A	5/1994 Gannon et al.			
5,319,760	A	6/1994 Mason et al.			
5,361,375	A	11/1994 Ogi			
5,386,552	A	1/1995 Garney			
5,421,006	A	5/1995 Jablon et al.			
5,434,999	A	7/1995 Goire et al.			
5,437,033	A	7/1995 Inoue et al.			
5,442,645	A	8/1995 Ugon et al.			
5,455,909	A	10/1995 Blomgren et al.			
5,459,867	A	10/1995 Adams et al.			
5,459,869	A	10/1995 Spilo			
5,469,557	A	11/1995 Salt et al.			
5,473,692	A	12/1995 Davis			
5,479,509	A	12/1995 Ugon			
5,504,922	A	4/1996 Seki et al.			
5,506,975	A	4/1996 Onodera			
5,511,217	A	4/1996 Nakajima et al.			
5,522,075	A	5/1996 Robinson et al.			
5,528,231	A	6/1996 Patarin			
5,533,126	A	7/1996 Hazard et al.			
5,555,385	A	9/1996 Osisek			
5,555,414	A	9/1996 Hough et al.			
5,560,013	A	9/1996 Scalzi et al.			
5,564,040	A	10/1996 Kubala			
5,566,323	A	10/1996 Ugon			
5,568,552	A	10/1996 Davis			
5,574,936	A	11/1996 Ryba et al.			
5,582,717	A	12/1996 Di Santo			
5,604,805	A	2/1997 Brands			
5,606,617	A	2/1997 Brands			
5,615,263	A	3/1997 Takahashi			
5,628,022	A	5/1997 Ueno et al.			
5,633,929	A	5/1997 Kaliski, Jr.			
5,657,445	A	8/1997 Pearce			
5,668,971	A	9/1997 Neufeld			
5,684,948	A	11/1997 Johnson et al.			
5,706,469	A	1/1998 Kobayashi			
5,717,903	A	2/1998 Bonola			
5,720,609	A	2/1998 Pfeifferle			
5,721,222	A	2/1998 Bernstein et al.			
5,729,760	A	3/1998 Poisner			
5,737,604	A	4/1998 Miller et al.			
5,737,760	A	4/1998 Grimmer, Jr. et al.			
5,740,178	A	4/1998 Jacks et al.			
5,752,046	A	5/1998 Oprescu et al.			
5,757,919	A	5/1998 Herbert et al.			

6,339,815	B1	1/2002	Feng	
6,339,816	B1	1/2002	Bausch	
6,357,004	B1	3/2002	Davis	
6,363,485	B1	3/2002	Adams	
6,374,286	B1	4/2002	Gee et al.	
6,374,317	B1	4/2002	Ajanovic et al.	
6,378,068	B1	4/2002	Foster	
6,378,072	B1	4/2002	Collins et al.	
6,389,537	B1	5/2002	Davis et al.	
6,397,242	B1	5/2002	Devine et al.	
6,397,379	B1	5/2002	Yates, Jr. et al.	
6,412,035	B1	6/2002	Webber	
6,421,702	B1	7/2002	Gulick	
6,435,416	B1	8/2002	Slassi	
6,445,797	B1	9/2002	McGough et al.	
6,463,535	B1	10/2002	Drews et al.	
6,463,537	B1	10/2002	Tello	
6,499,123	B1	12/2002	McFarland et al.	
6,505,279	B1	1/2003	Phillips et al.	
6,507,904	B1	1/2003	Ellison et al.	
6,529,909	B1	3/2003	Bowman-Amuah	
6,535,988	B1	3/2003	Poisner	
6,557,104	B2	4/2003	Vu et al.	
6,560,627	B1	5/2003	McDonald et al.	
6,609,199	B1	8/2003	DeTreville	
6,615,278	B1	9/2003	Curtis	
6,633,963	B1	10/2003	Ellison et al.	
6,633,981	B1	10/2003	Davis	
6,651,171	B1	11/2003	England et al.	
6,678,825	B1	1/2004	Ellison et al.	
6,684,326	B1	1/2004	Cromer et al.	
6,961,806	B1 *	11/2005	Agesen et al.	711/6
2001/0021969	A1	9/2001	Burger et al.	
2001/0027511	A1	10/2001	Wakabayashi et al.	
2001/0027527	A1	10/2001	Khidekel et al.	
2001/0037450	A1	11/2001	Metlitski et al.	
2002/0007456	A1	1/2002	Peinado et al.	
2002/0023032	A1	2/2002	Pearson et al.	
2002/0082824	A1	6/2002	Neiger et al.	
2002/0147916	A1	10/2002	Strongin et al.	
2002/0166061	A1	11/2002	Falik et al.	
2002/0169717	A1	11/2002	Challener	
2003/0018892	A1	1/2003	Tello	
2003/0074548	A1	4/2003	Cromer et al.	
2003/0115453	A1	6/2003	Grawrock	
2003/0126442	A1	7/2003	Glew et al.	
2003/0126453	A1	7/2003	Glew et al.	
2003/0159056	A1	8/2003	Cromer et al.	
2003/0188179	A1	10/2003	Challener et al.	
2003/0196085	A1	10/2003	Lampson et al.	
2004/0064668	A1	4/2004	Kjos et al.	
2004/0117539	A1	6/2004	Bennett et al.	
2006/0026384	A1 *	2/2006	Brandt et al.	711/209

FOREIGN PATENT DOCUMENTS

EP	0473913	3/1992
EP	0600112	6/1994
EP	0602867	6/1994
EP	0892521	1/1999
EP	0930567	7/1999
EP	0961193	12/1999
EP	0965902	12/1999
EP	1030237	8/2000
EP	1055989	11/2000
EP	1056014	11/2000
EP	1085396	3/2001
EP	1146715	10/2001
EP	1209563	5/2002
EP	1271277	1/2003
JP	2000076139	3/2000
WO	WO9524696	9/1995
WO	WO9729567	8/1997

WO	WO9812620	3/1998
WO	WO9834365	8/1998
WO	WO9844402	10/1998
WO	WO9905600	2/1999
WO	WO9918511	4/1999
WO	WO9957863	11/1999
WO	WO9965579	12/1999
WO	WO0021238	4/2000
WO	WO0062232	10/2000
WO	WO0127723	4/2001
WO	WO0127821	4/2001
WO	WO0163994	8/2001
WO	WO0175565	10/2001
WO	WO0175595	10/2001
WO	WO0201794	1/2002
WO	WO9909482	1/2002
WO	WO0217555	2/2002
WO	WO02060121	8/2002
WO	WO0175564	10/2002
WO	WO02086684	10/2002
WO	WO03058412	7/2003

OTHER PUBLICATIONS

Brands, Stefan, "Restrictive Blinding of Secret-Key Certificates", *Springer-Verlag XP002201306*, (1995), Chapter 3.

Chien, Andrew A., et al., "Safe and Protected Execution for the Morph/AMRM Reconfigurable Processor", *7th Annual IEEE Symposium, FCCM '99 Proceedings*, XP010359180, ISBN 0-7695-0375-6, Los Alamitos, CA, (Apr. 21, 1999), 209-221.

Compaq Computer Corporation, "Trusted Computing Platform Alliance (TCPA) Main Specification Version 1.1a", *XP002272822*, (Jan. 25, 2001), 1-321.

Coulouris, George, et al., "Distributed Systems, Concepts and Designs", *2nd Edition*, (1994), 422-424.

Crawford, John, "Architecture of the Intel 80386", *Proceedings of the IEEE International Conference on Computer Design: VLSI in Computers and Processors (ICCD '86)*, (Oct. 6, 1986), 155-160.

David, George I., et al., "Defending Systems Against Viruses through Cryptographic Authentication", *Proceedings of the Symposium on Security and Privacy*, IEEE Comp. Soc. Press, ISBN 0-8186-1939-2, (May 1989).

Fabry, R.S., "Capability-Based Addressing", *Communications of the ACM*, vol. 17, No. 7, (Jul. 1974), 403-412.

Frieder, Gideon, "The Architecture And Operational Characteristics of the VMX Host Machine", *The Architecture And Operational Characteristics of the VMX Host Machine, IEEE*, (1982), 9-16.

Goldberg, Robert P., "Survey of Virtual Machine Research", *Computer Magazine*, (Jun. 1974), 34-35.

Gong, Li, et al., "Going Beyond the Sandbox: An Overview of the New Security Architecture in the Java Development Kit 1.2", *Proceedings of the USENIX Symposium on Internet Technologies and Systems*, Monterey, CA, (Dec. 1997).

Gum, P. H., "System/370 Extended Architecture: Facilities for Virtual Machines", *IBM J. Research Development*, vol. 27, No. 6, (Nov. 1983), 530-544.

Heinrich, Joe, "MIPS R4000 Microprocessor User's Manual, Second Edition", *Chapter 4 "Memory Management"*, (Jun. 11, 1993), 61-97.

HP Mobile Security Overview, "HP Mobile Security Overview", (Sep. 2002), 1-10.

IBM, "Information Display Technique for a Terminate Stay Resident Program IBM Technical Disclosure Bulletin", *TDB-ACC-No. NA9112156*, vol. 34, Issue 7A, (Dec. 1, 1991), 156-158.

IBM Corporation, "IBM ThinkPad T30 Notebooks", *IBM Product Specification*, located at www-1.ibm.com/services/files/cisco_t30_spec_sheet_070202.pdf, last visited Jun. 23, 2004, (Jul. 2, 2002), 1-6.

Intel, "IA-32 Intel Architecture Software Developer's Manual", vol. 3: *System Programming Guide*, Intel Corporation—2003, 13-1 through 13-24.

Intel, "Intel386 DX Microprocessor 32-Bit CMOS Microprocessor With Integrated Memory Management", (1995), 5-56.

- Intel Corporation, "IA-64 System Abstraction Layer Specification", *Intel Product Specification*, Order No. 245359-001, (Jan. 2000), 1-112.
- Intel Corporation, "Intel 82802AB/82802AC Firmware Hub (FWH)", *Intel Product Datasheet*, Document No. 290658-004, (Nov. 2000), 1-6, 17-28.
- Intel Corporation, "Intel IA-64 Architecture Software Developer's Manual", vol. 2: *IA-64 System Architecture*, Order No. 245318-001, (Jan. 2000), i, ii, 5.1-5.3, 11.1-11.8, 11.23-11.26.
- Karger, Paul A., et al., "A VMM Security Kernel for the VAX Architecture", *Proceedings of the Symposium on Research in Security and Privacy*, XP010020182, ISBN 0-8186-2060-9, Boxborough, MA, (May 7, 1990), 2-19.
- Kashiwagi, Kazuhiko, et al., "Design and Implementation of Dynamically Reconstructing System Software", *Software Engineering Conference*, Proceedings 1996 Asia-Pacific Seoul, South Korea Dec. 4-7, 1996, Los Alamitos, CA USA, IEEE Comput. Soc, US, ISBN 0-8186-7638-8, (1996).
- Lawton, Kevin, et al., "Running Multiple Operating Systems Concurrently on an IA32 PC Using Virtualization Techniques", <http://www.plex86.org/research/paper.txt>, (Nov. 29, 1999), 1-31.
- Luke, Jahn, et al., "Replacement Strategy for Aging Avionics Computers", *IEEE AES Systems Magazine*, XP002190614, (Mar. 1999).
- Menezes, Alfred J., et al., "Handbook of Applied Cryptography", *CRC Press Series on Discrete Mathematics and its Applications*, Boca Raton, FL, XP002165287, ISBN 0849385237, (Oct. 1996), 403-405, 506-515, 570.
- Menezes, Alfred J., et al., "Handbook of Applied Cryptography", *CRC Press LLC*, USA XP002201307, 1997), 475.
- Motorola, "M68040 User's Manual", (1993), 1-1 to 8-32.
- Nanba, S., et al., "VM/4: ACOS-4 Virtual Machine Architecture", *VM/4: ACOS-4 Virtual Machine Architecture*, *IEEE*, (1985), 171-178.
- Richt, Stefan, et al., "In-Circuit-Emulator Wird Echtzeittaughlich", *Elektronik*, Franzis Verlag GMBH, Munchen, DE, vol. 40, No. 16, XP000259620, (100-103), Aug. 6, 1991.
- Robin, John S., et al., "Analysis of the Pentium's Ability to Support a Secure Virtual Machine Monitor", *Proceedings of the 9th USENIX Security Symposium*, XP002247347, Denver, Colorado, (Aug. 14, 2000), 1-17.
- Rosenblum, M., "Virtual Platform: A Virtual Machine Monitor for Commodity PC", *Proceedings of the 11th Hotchips Conference*, (Aug. 17, 1999), 185-196.
- RSA Security, "Hardware Authenticators", www.rsasecurity.com/node.asp?id=1158, 1-2.
- RSA Security, "RSA SecurID Authenticators", www.rsasecurity.com/products/securid/datasheets/SID_DS_0103.pdf, 1-2.
- RSA Security, "Software Authenticators", www.rsasecurity.com/node.asp?id=1313, 1-2.
- Saez, Sergio, et al., "A Hardware Scheduler for Complex Real-Time Systems", *Proceedings of the IEEE International Symposium on Industrial Electronics*, XP002190615, (Jul. 1999), 43-48.
- Schneier, Bruce, "Applied Cryptography: Protocols, Algorithm, and Source Code in C", *Wiley, John & Sons, Inc.*, XP002939871; ISBN 0471117099, (Oct. 1995), 47-52.
- Schneier, Bruce, "Applied Cryptography: Protocols, Algorithm, and Source Code in C", *Wiley, John & Sons, Inc.*, XP002138607; ISBN 0471117099, (Oct. 1995), 56-65.
- Schneier, Bruce, "Applied Cryptography: Protocols, Algorithms, and Source Code C", *Wiley, John & Sons, Inc.*, XP002111449; ISBN 0471117099, (Oct. 1995), 169-187.
- Schneier, Bruce, "Applied Cryptography: Protocols, Algorithms, and Source Code in C", *2nd Edition; Wiley, John & Sons, Inc.*, XP002251738; ISBN 0471128457, (Nov. 1995), 28-33; 176-177; 216-217; 461-473; 518-522.
- Sherwood, Timothy, et al., "Patchable Instruction ROM Architecture", *Department of Computer Science and Engineering*, University of California, San Diego, La Jolla, CA, (Nov. 2001).
- Hall, Judith S., et al., "Virtualizing the VAX Architecture," *ACM SIGARCH Computer Architecture News*, Proceedings of the 18th annual international symposium on Computer architecture, vol. 19, Issue No. 3, Apr. 1991, 10 pages.
- Barham, Paul, et al., "Xen and the Art of Virtualization," *Proceedings of the ACM Symposium on Operating Systems Principles*, Oct. 2003, pp. 164-177.
- PCT Search Report PCT/US2006/003587, mailed Jul. 19, 2006, 6 pages.

* cited by examiner






FIG. 1

200



"Virtual TLB" 204



PD = page directory
 PT = page table
 F = page frame

FIG. 2




FIG. 3A

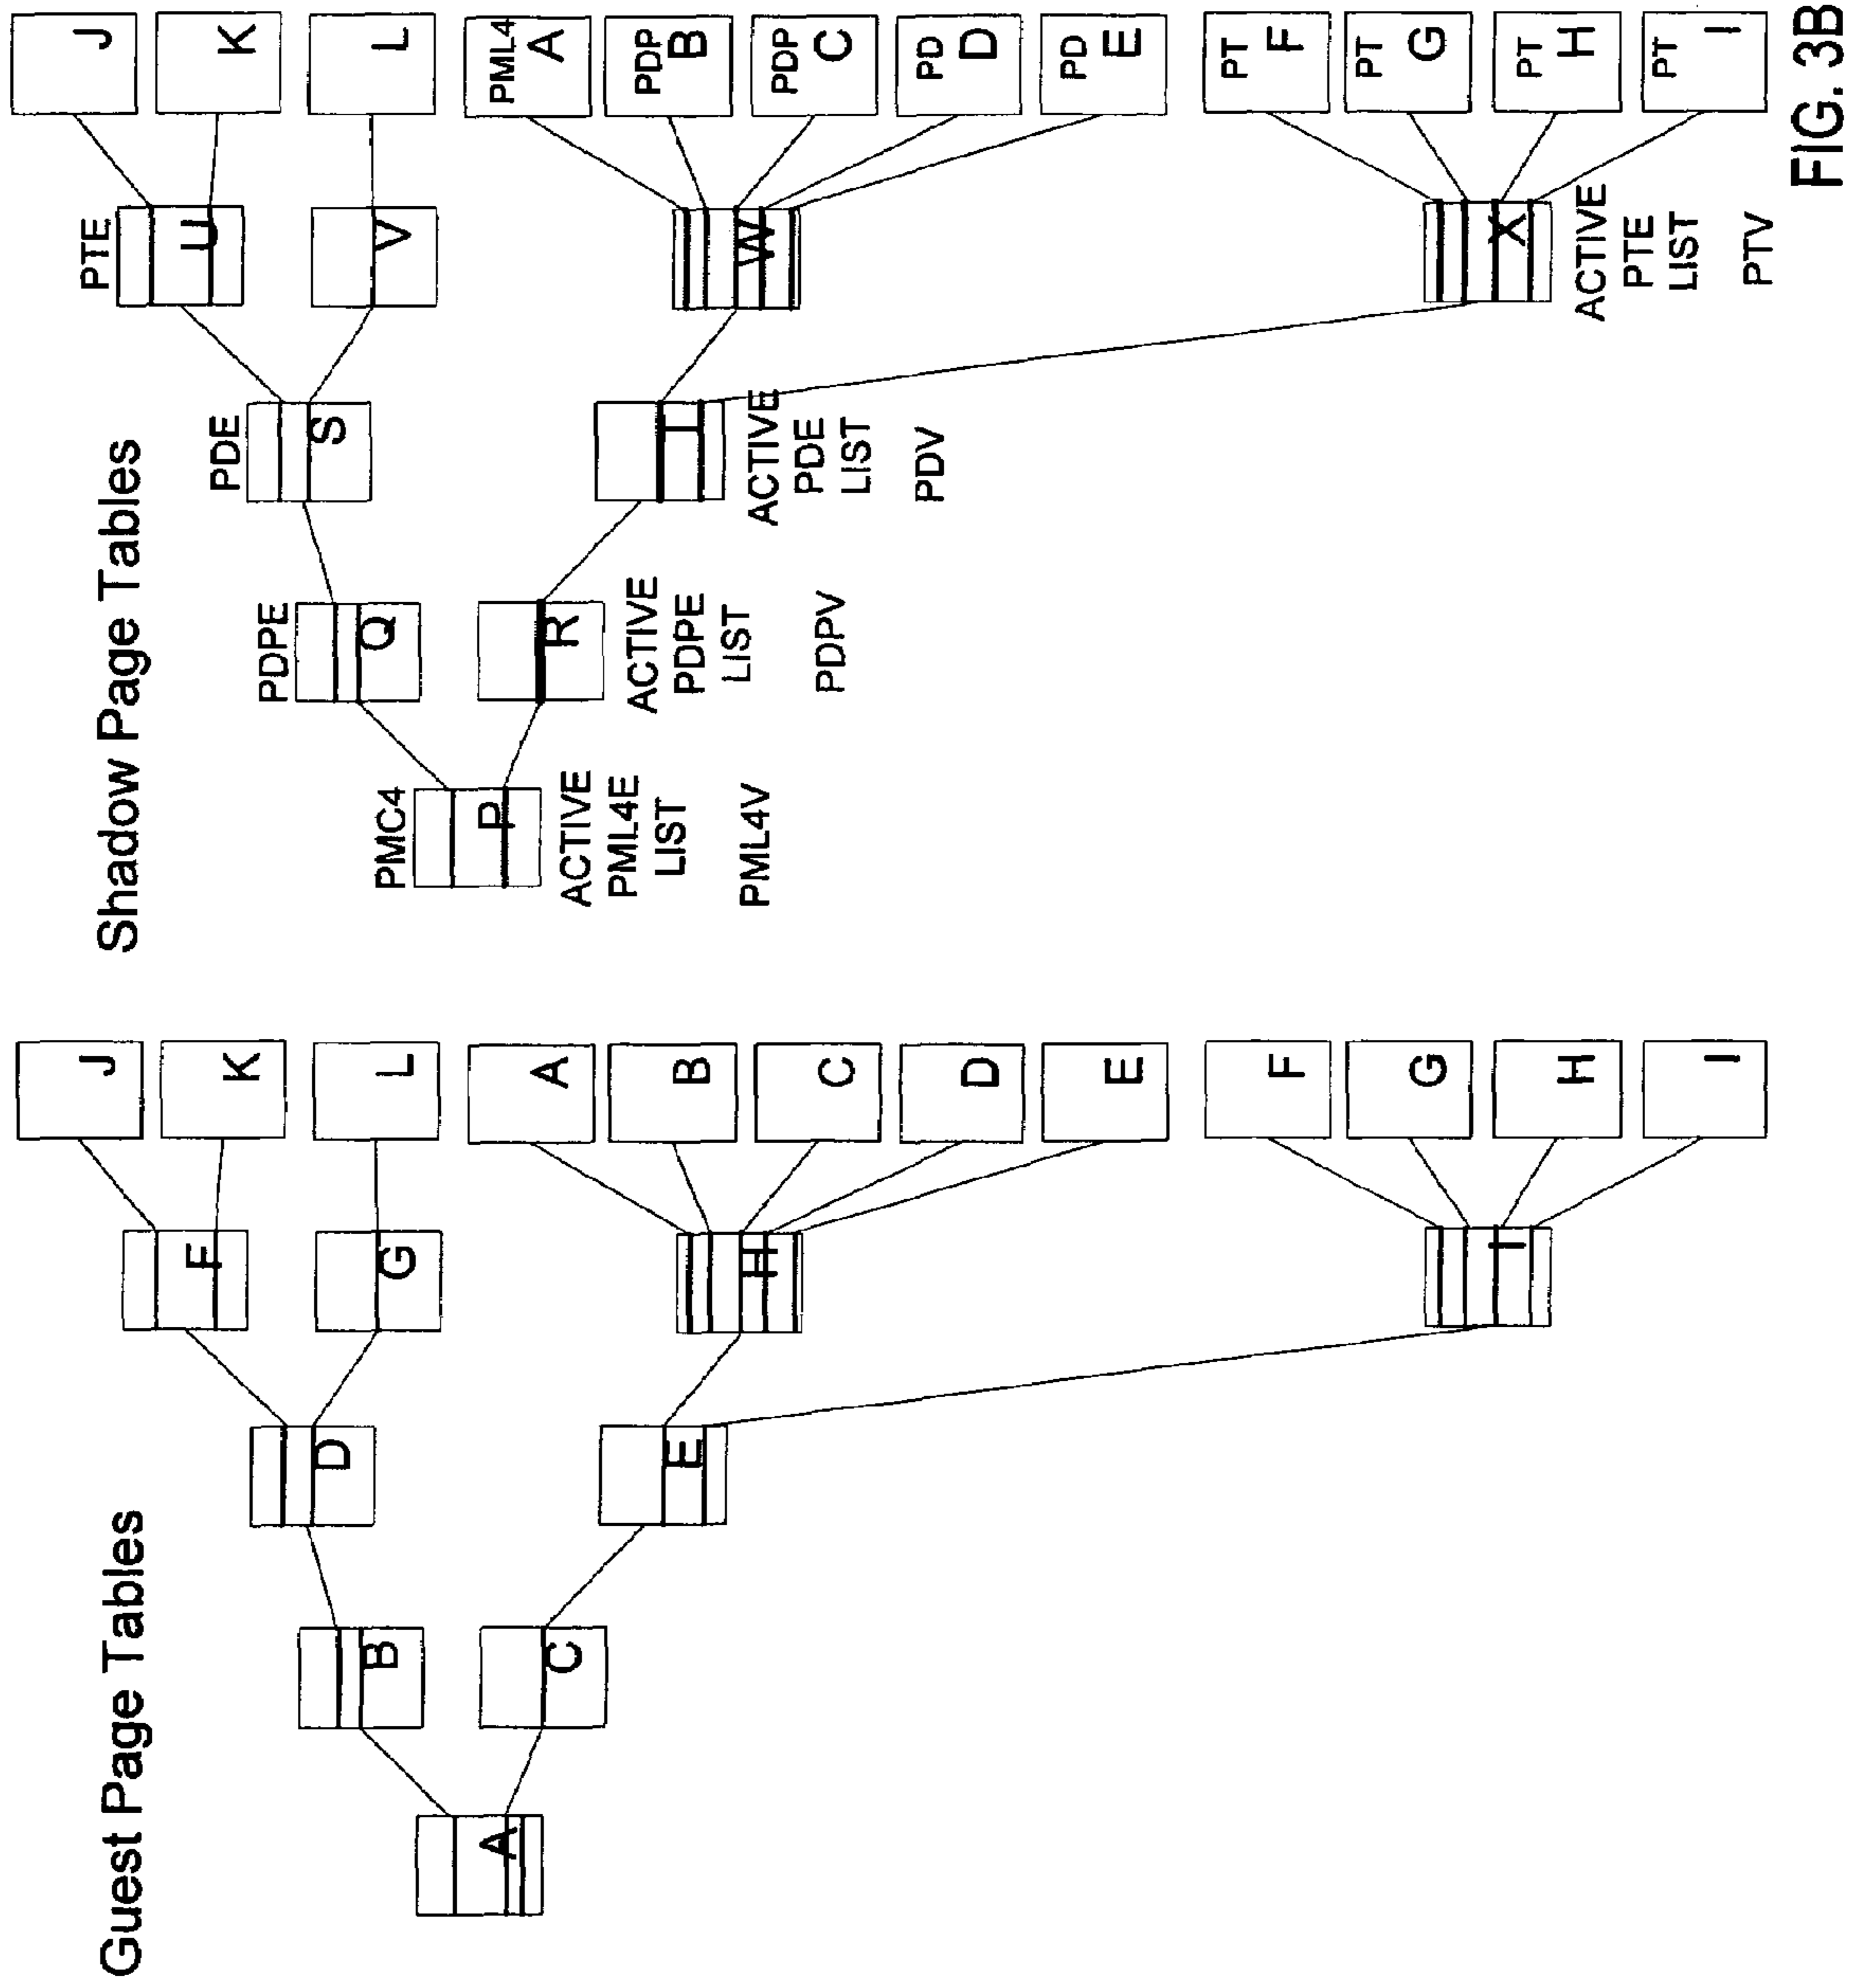


FIG. 3B




FIG. 4




FIG. 5




FIG. 6

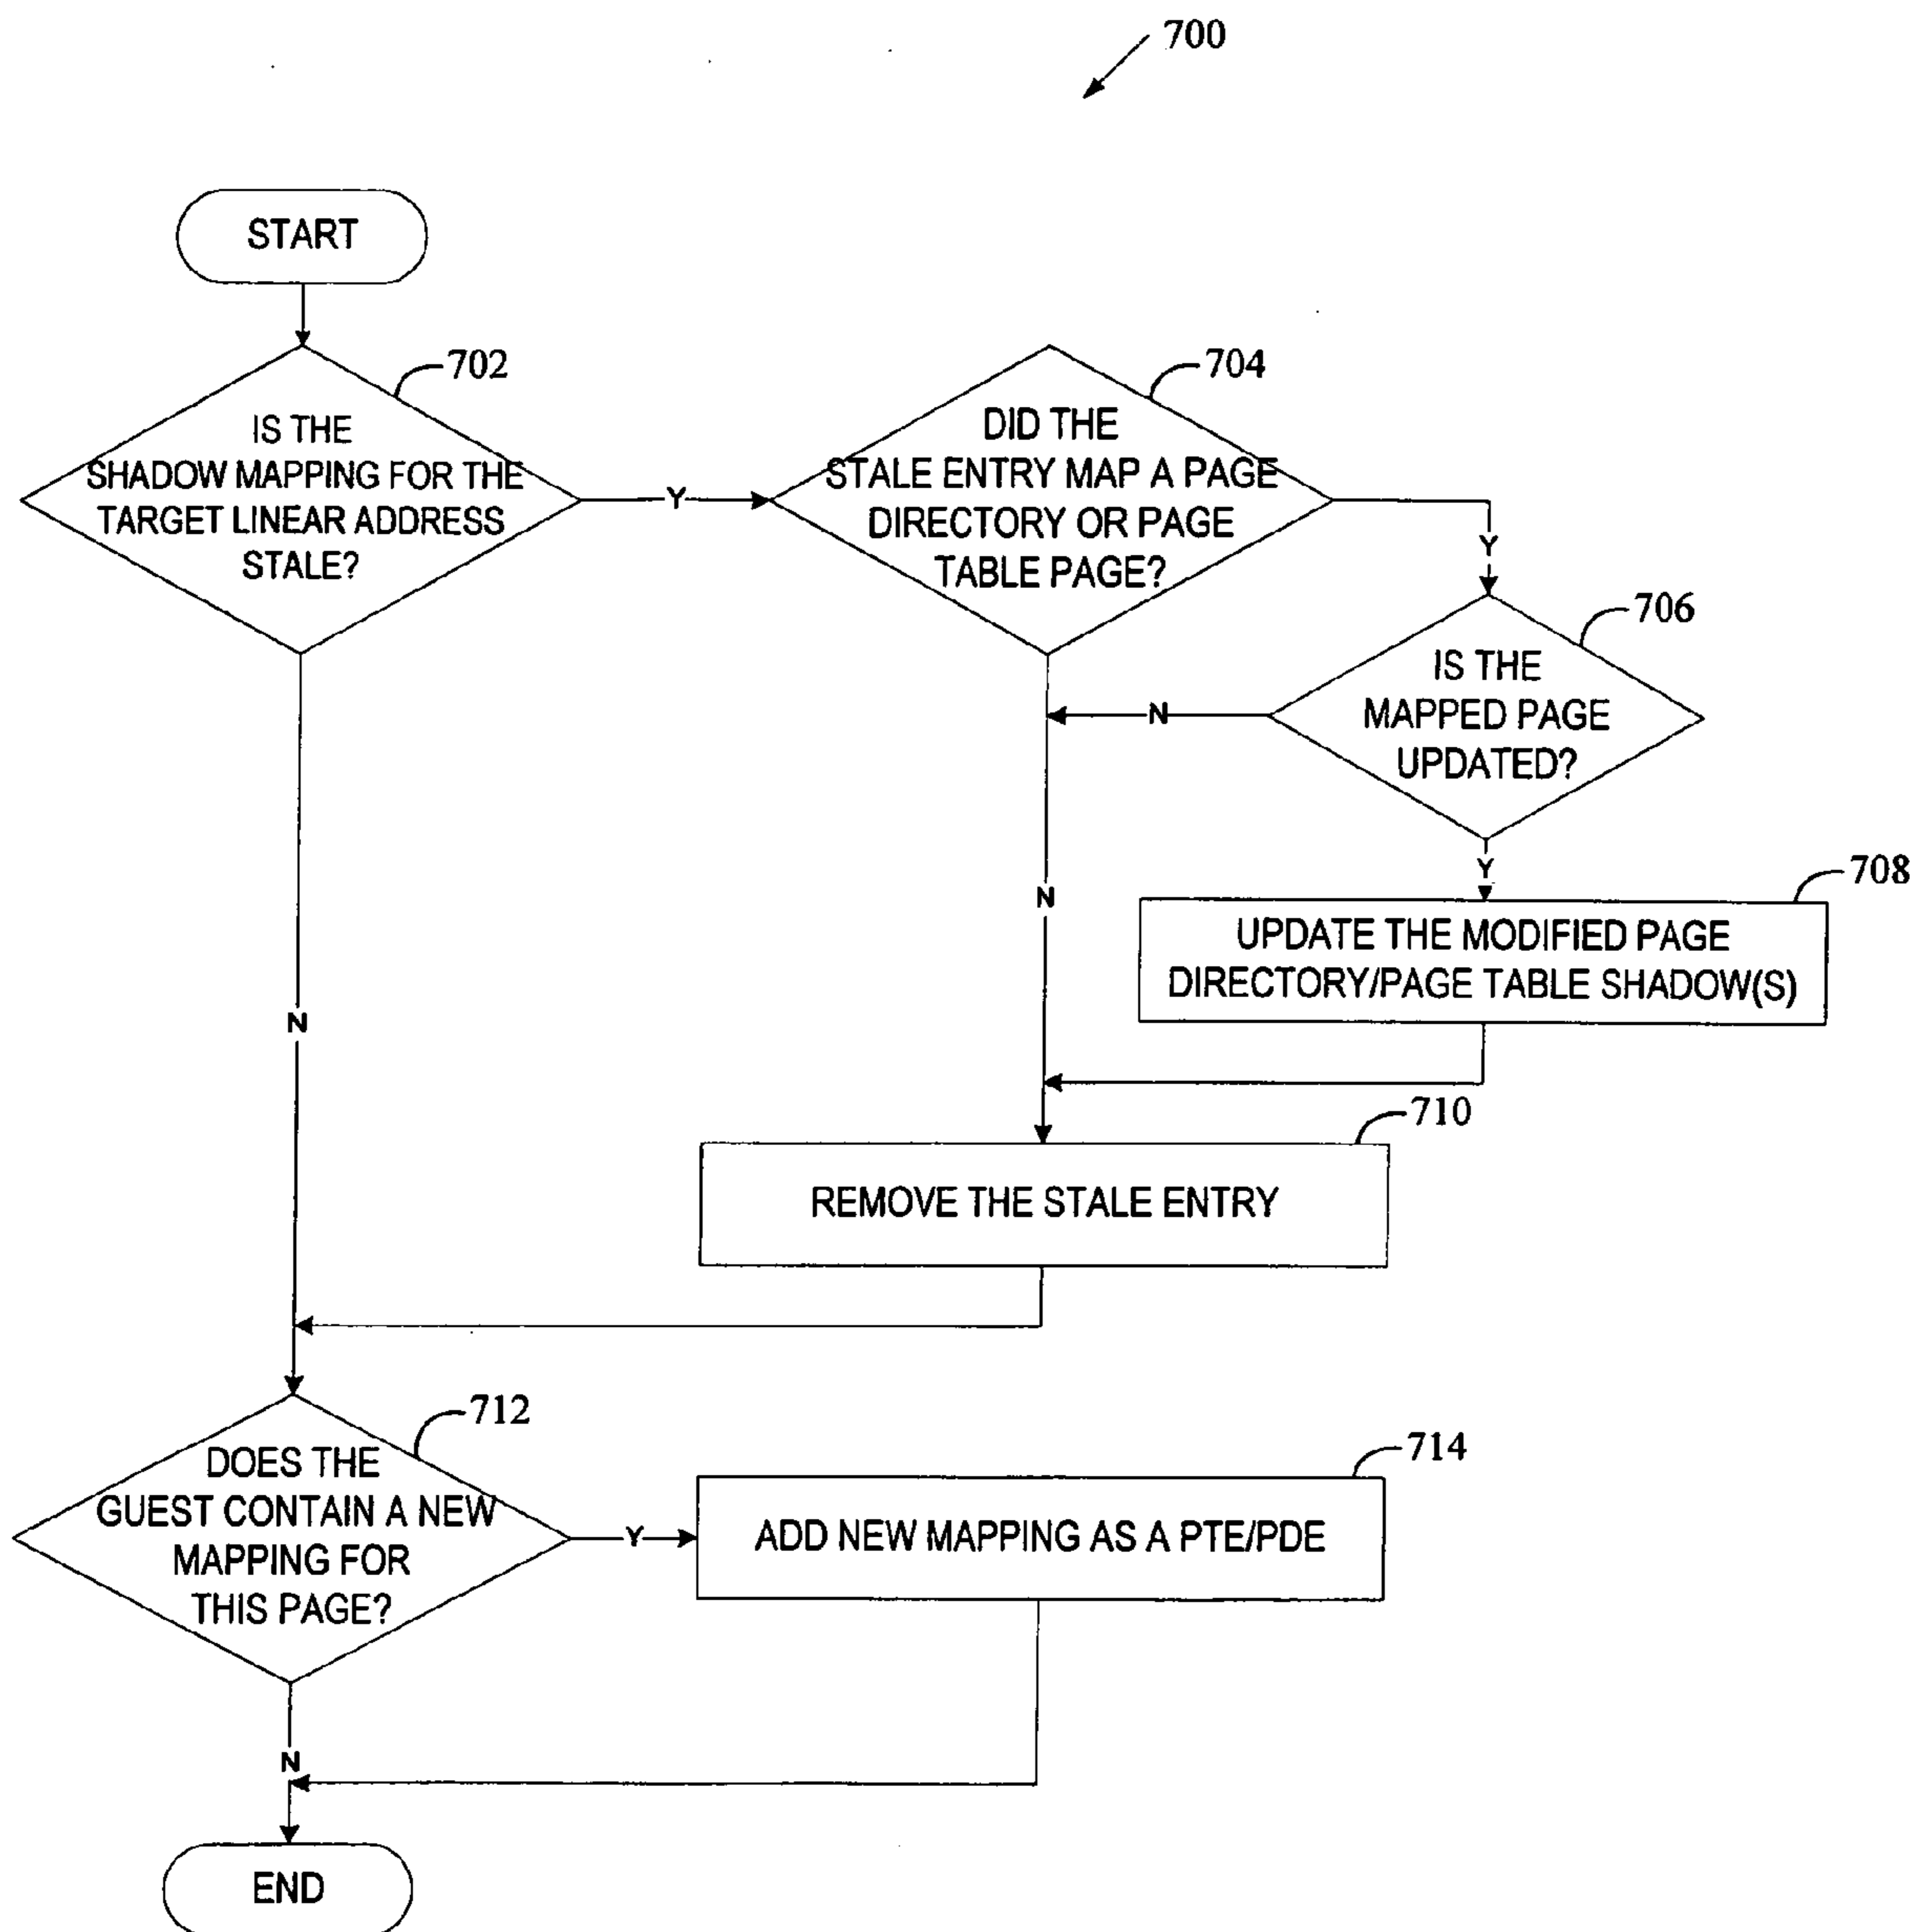


FIG. 7

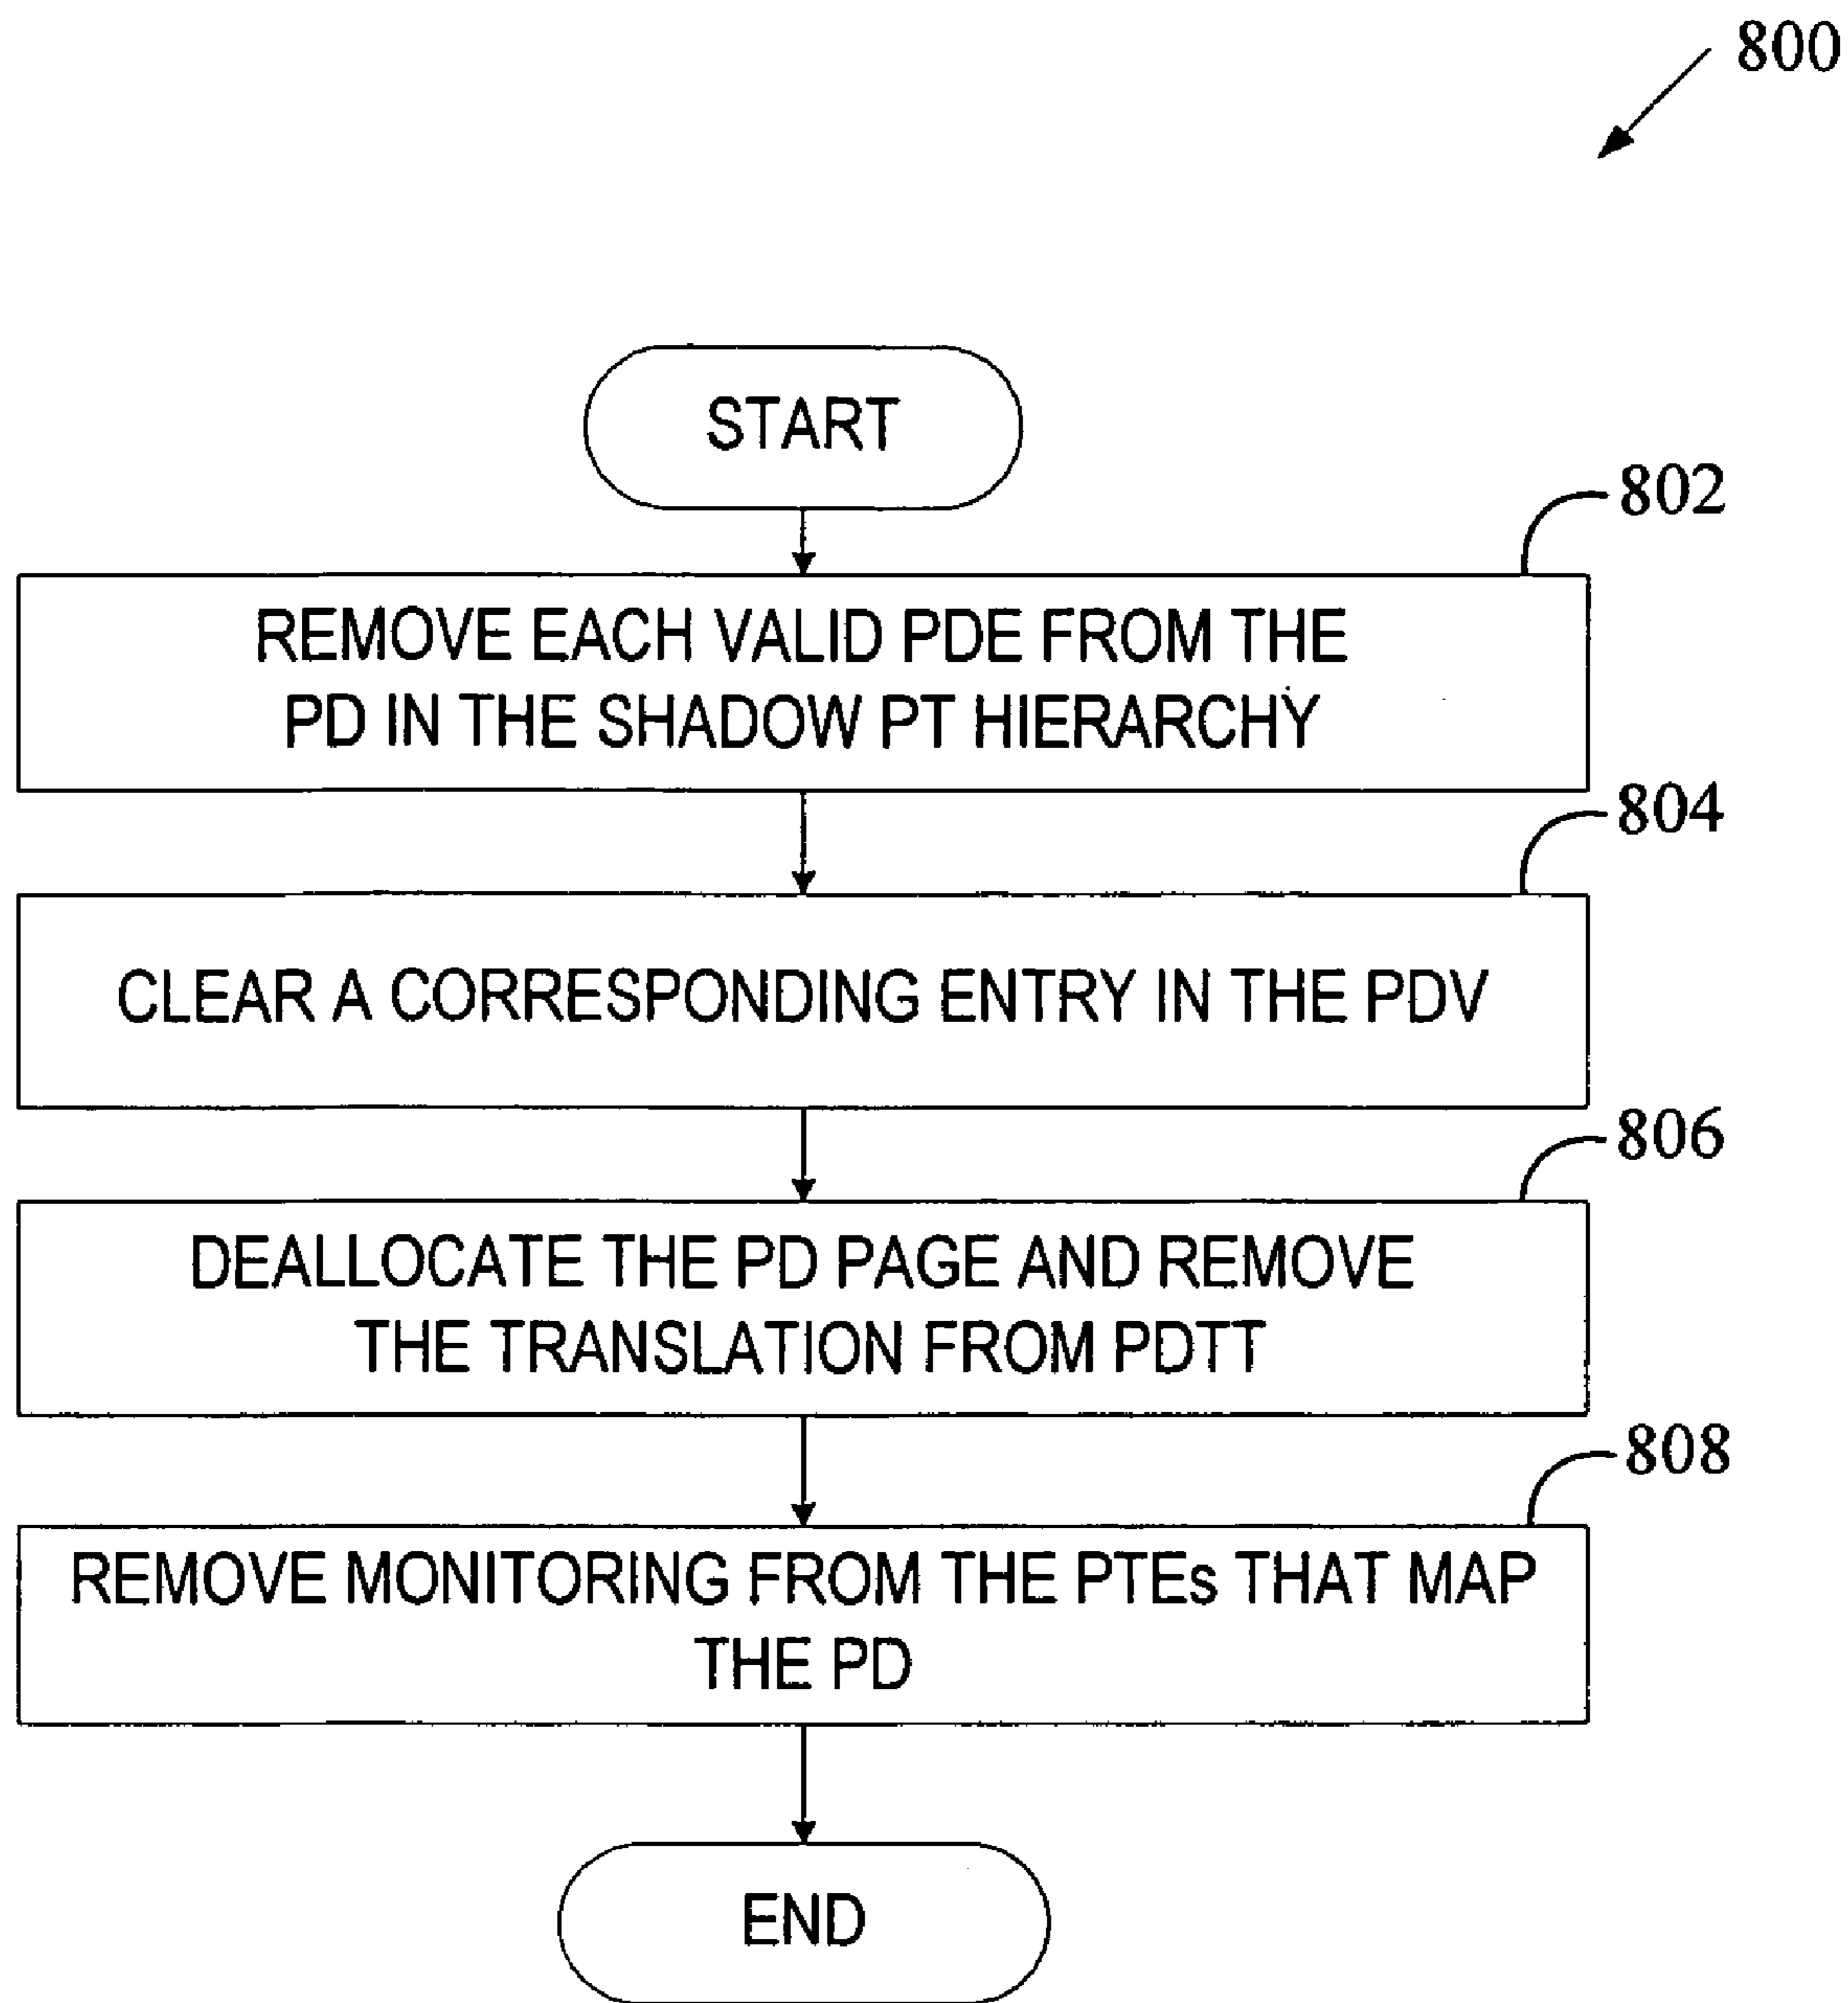


FIG. 8

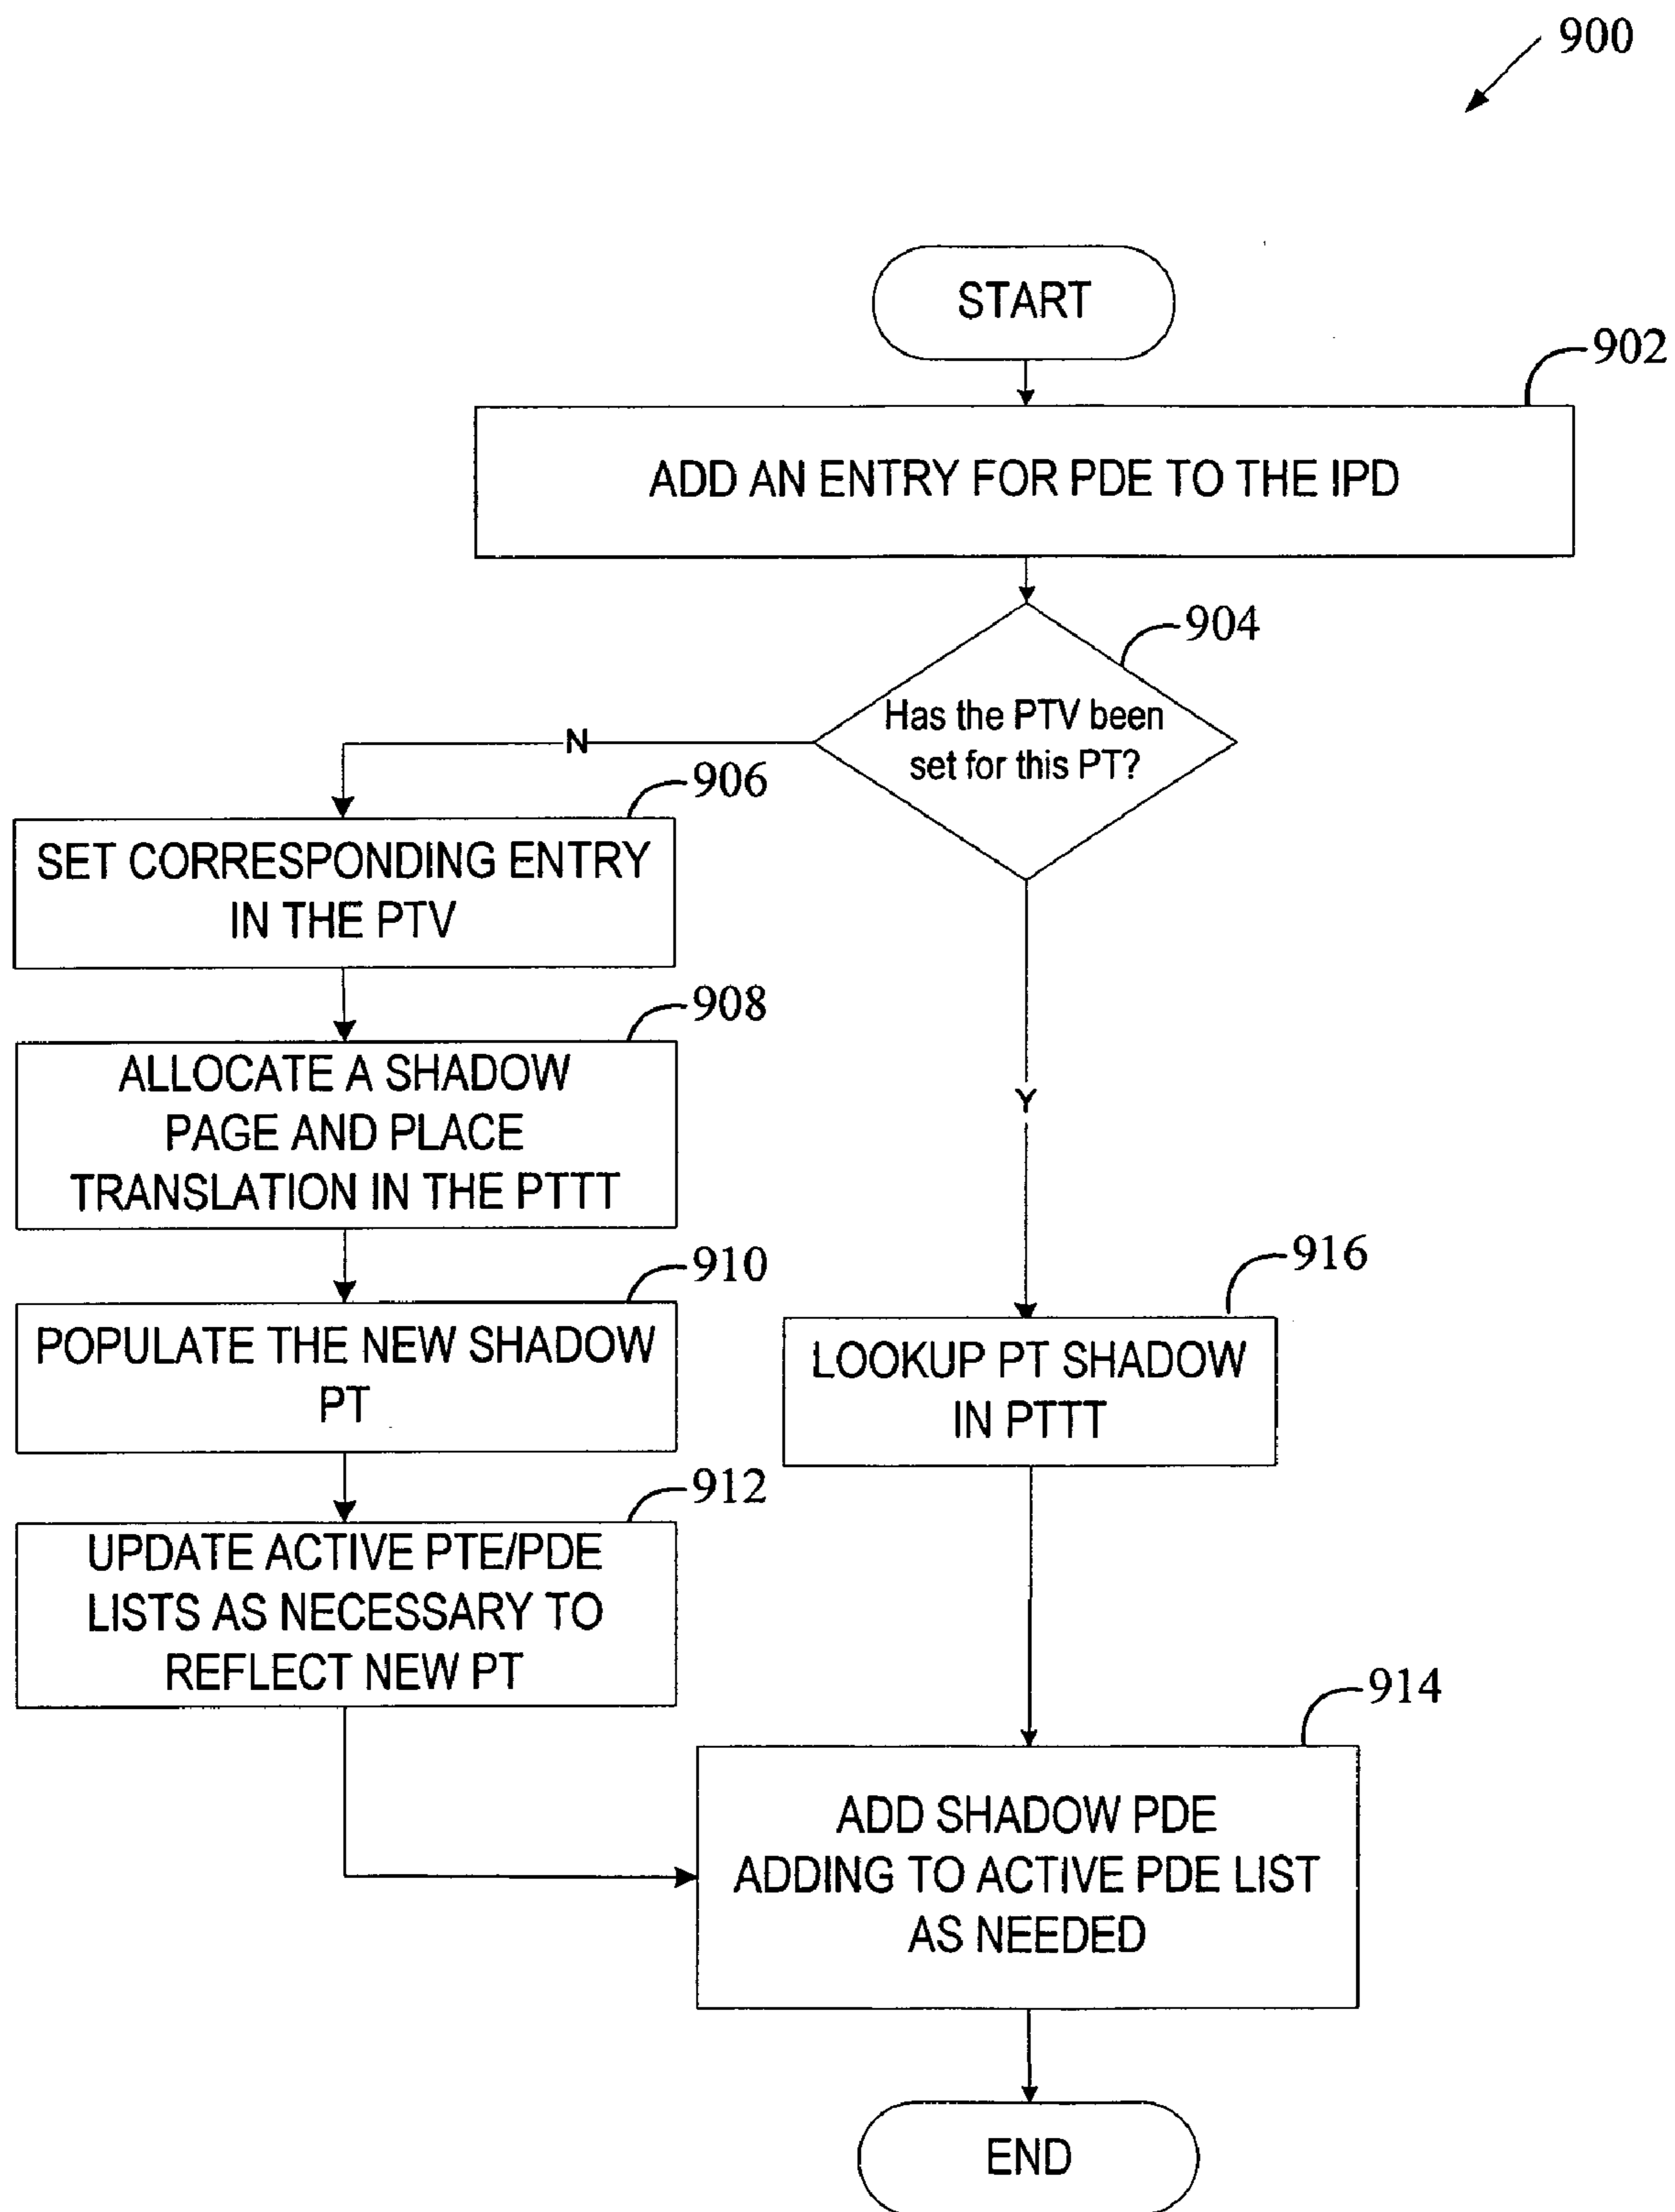


FIG. 9

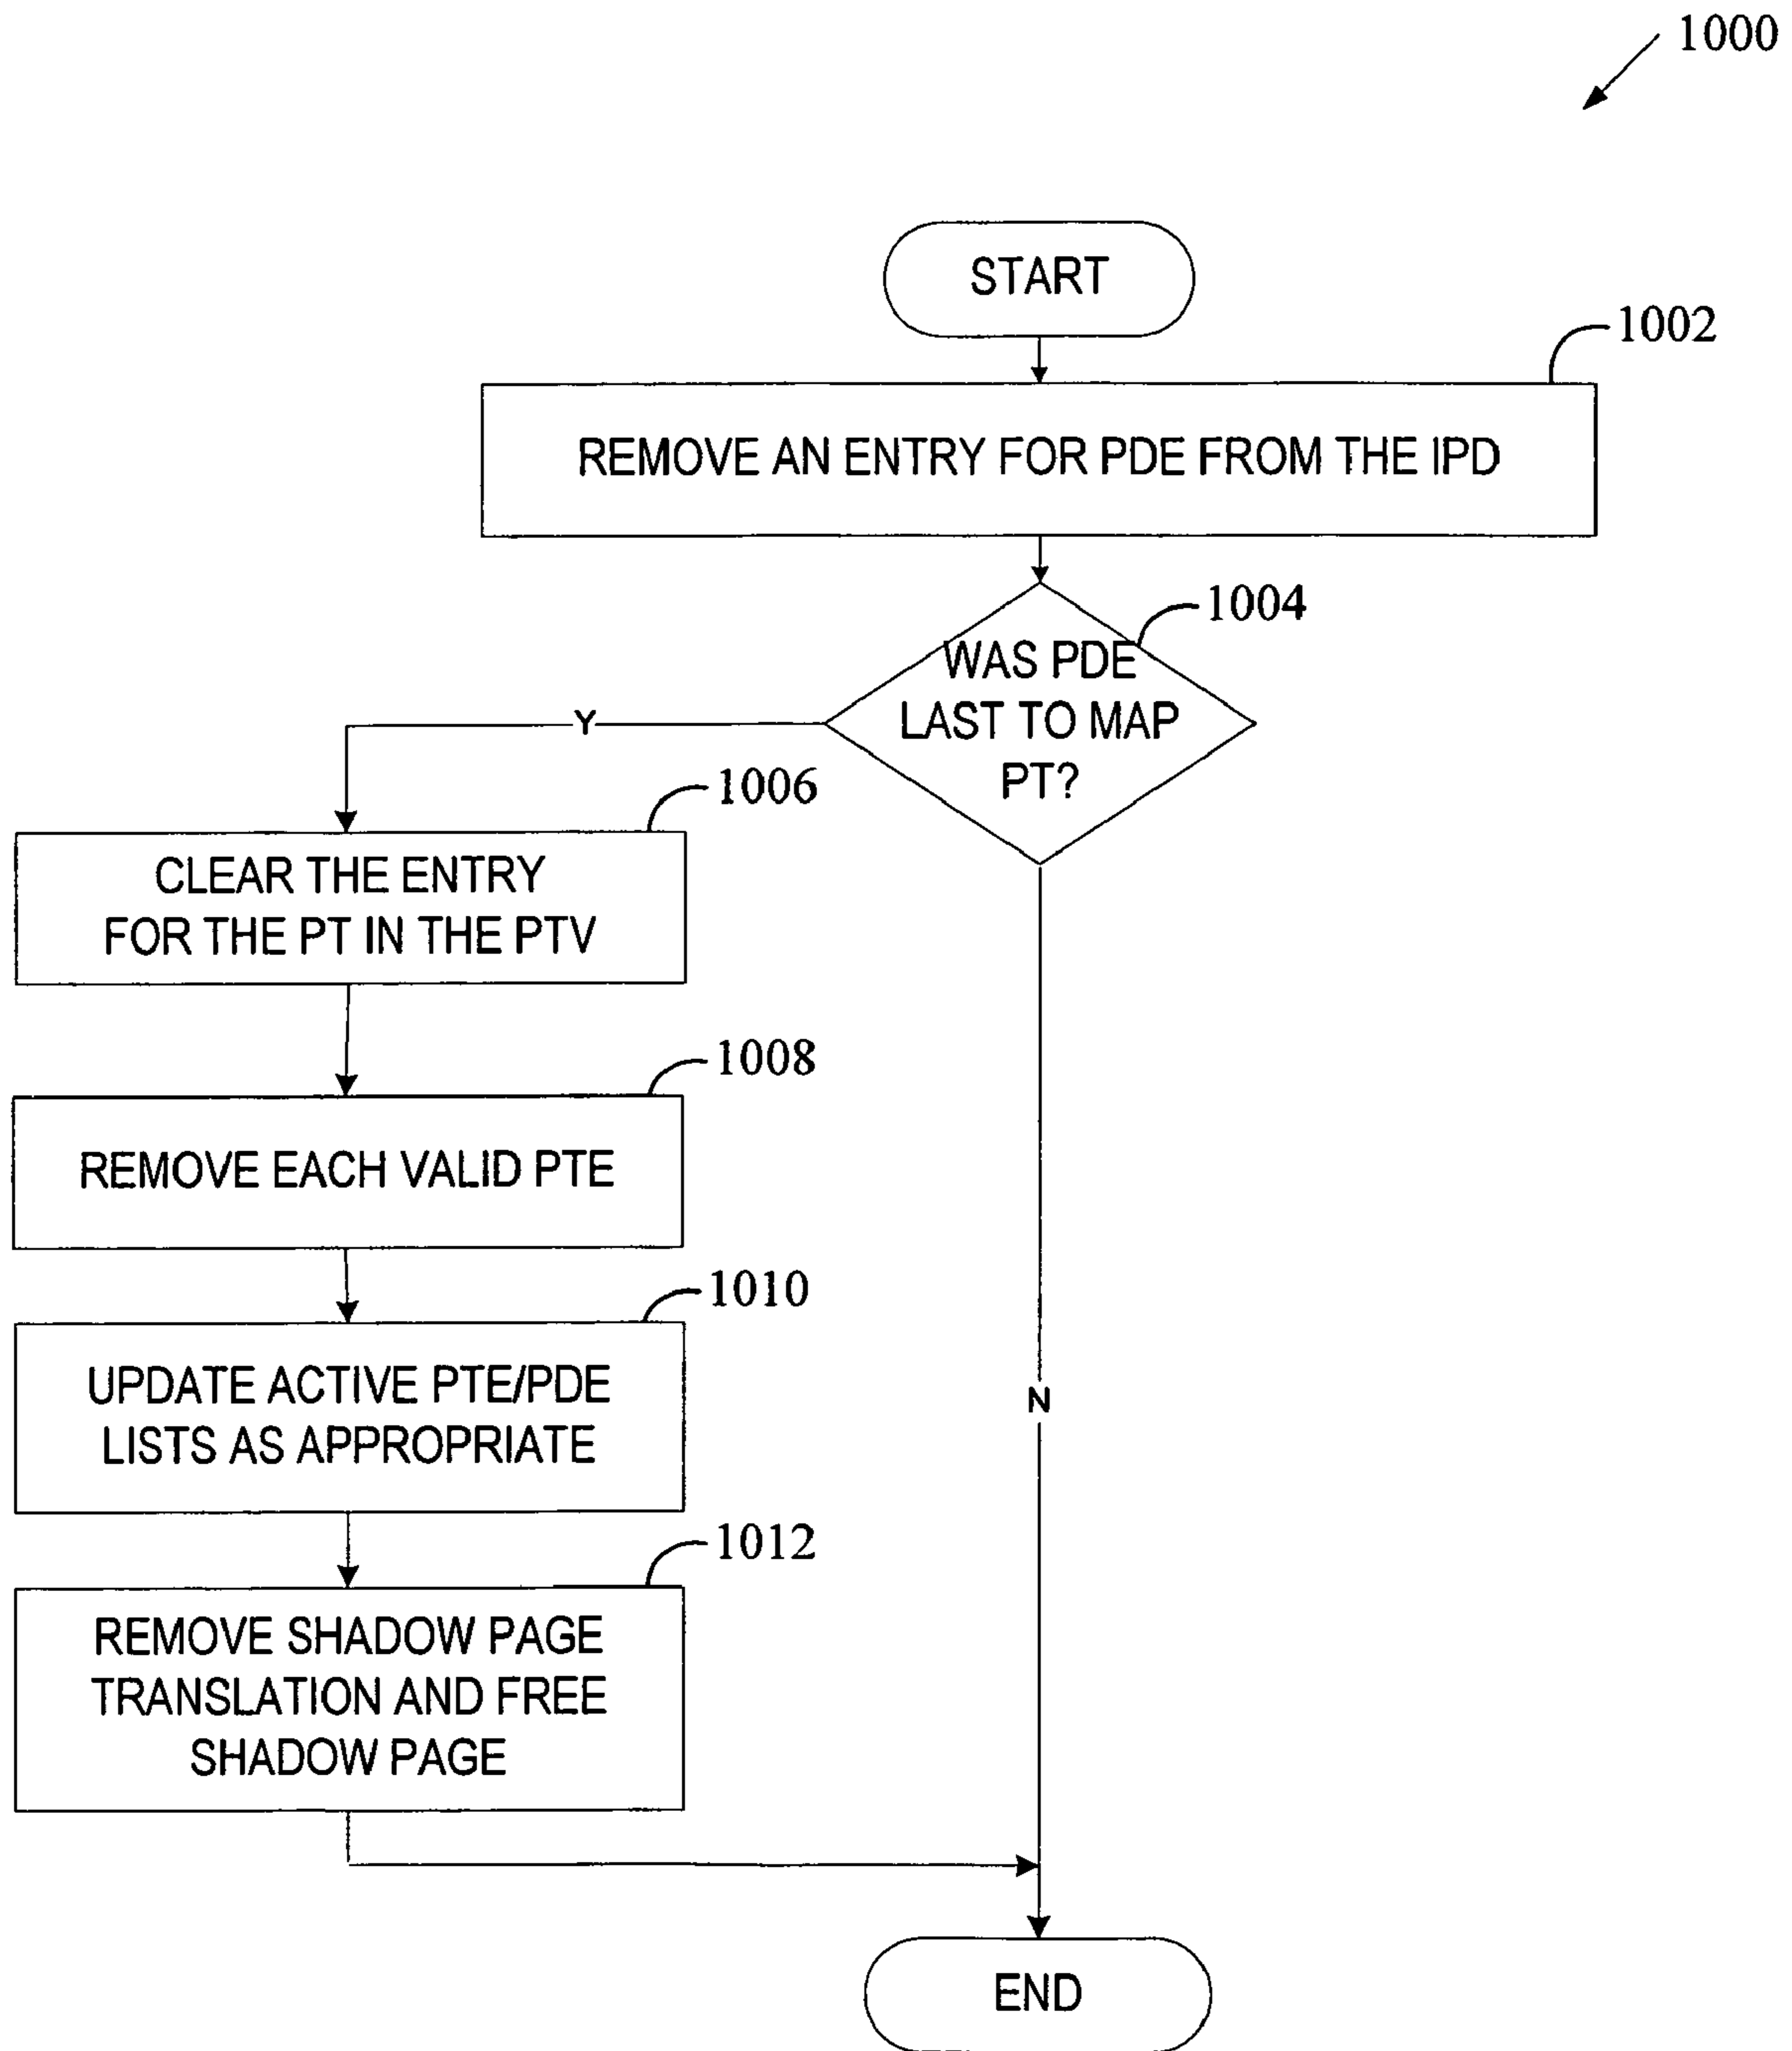


FIG. 10

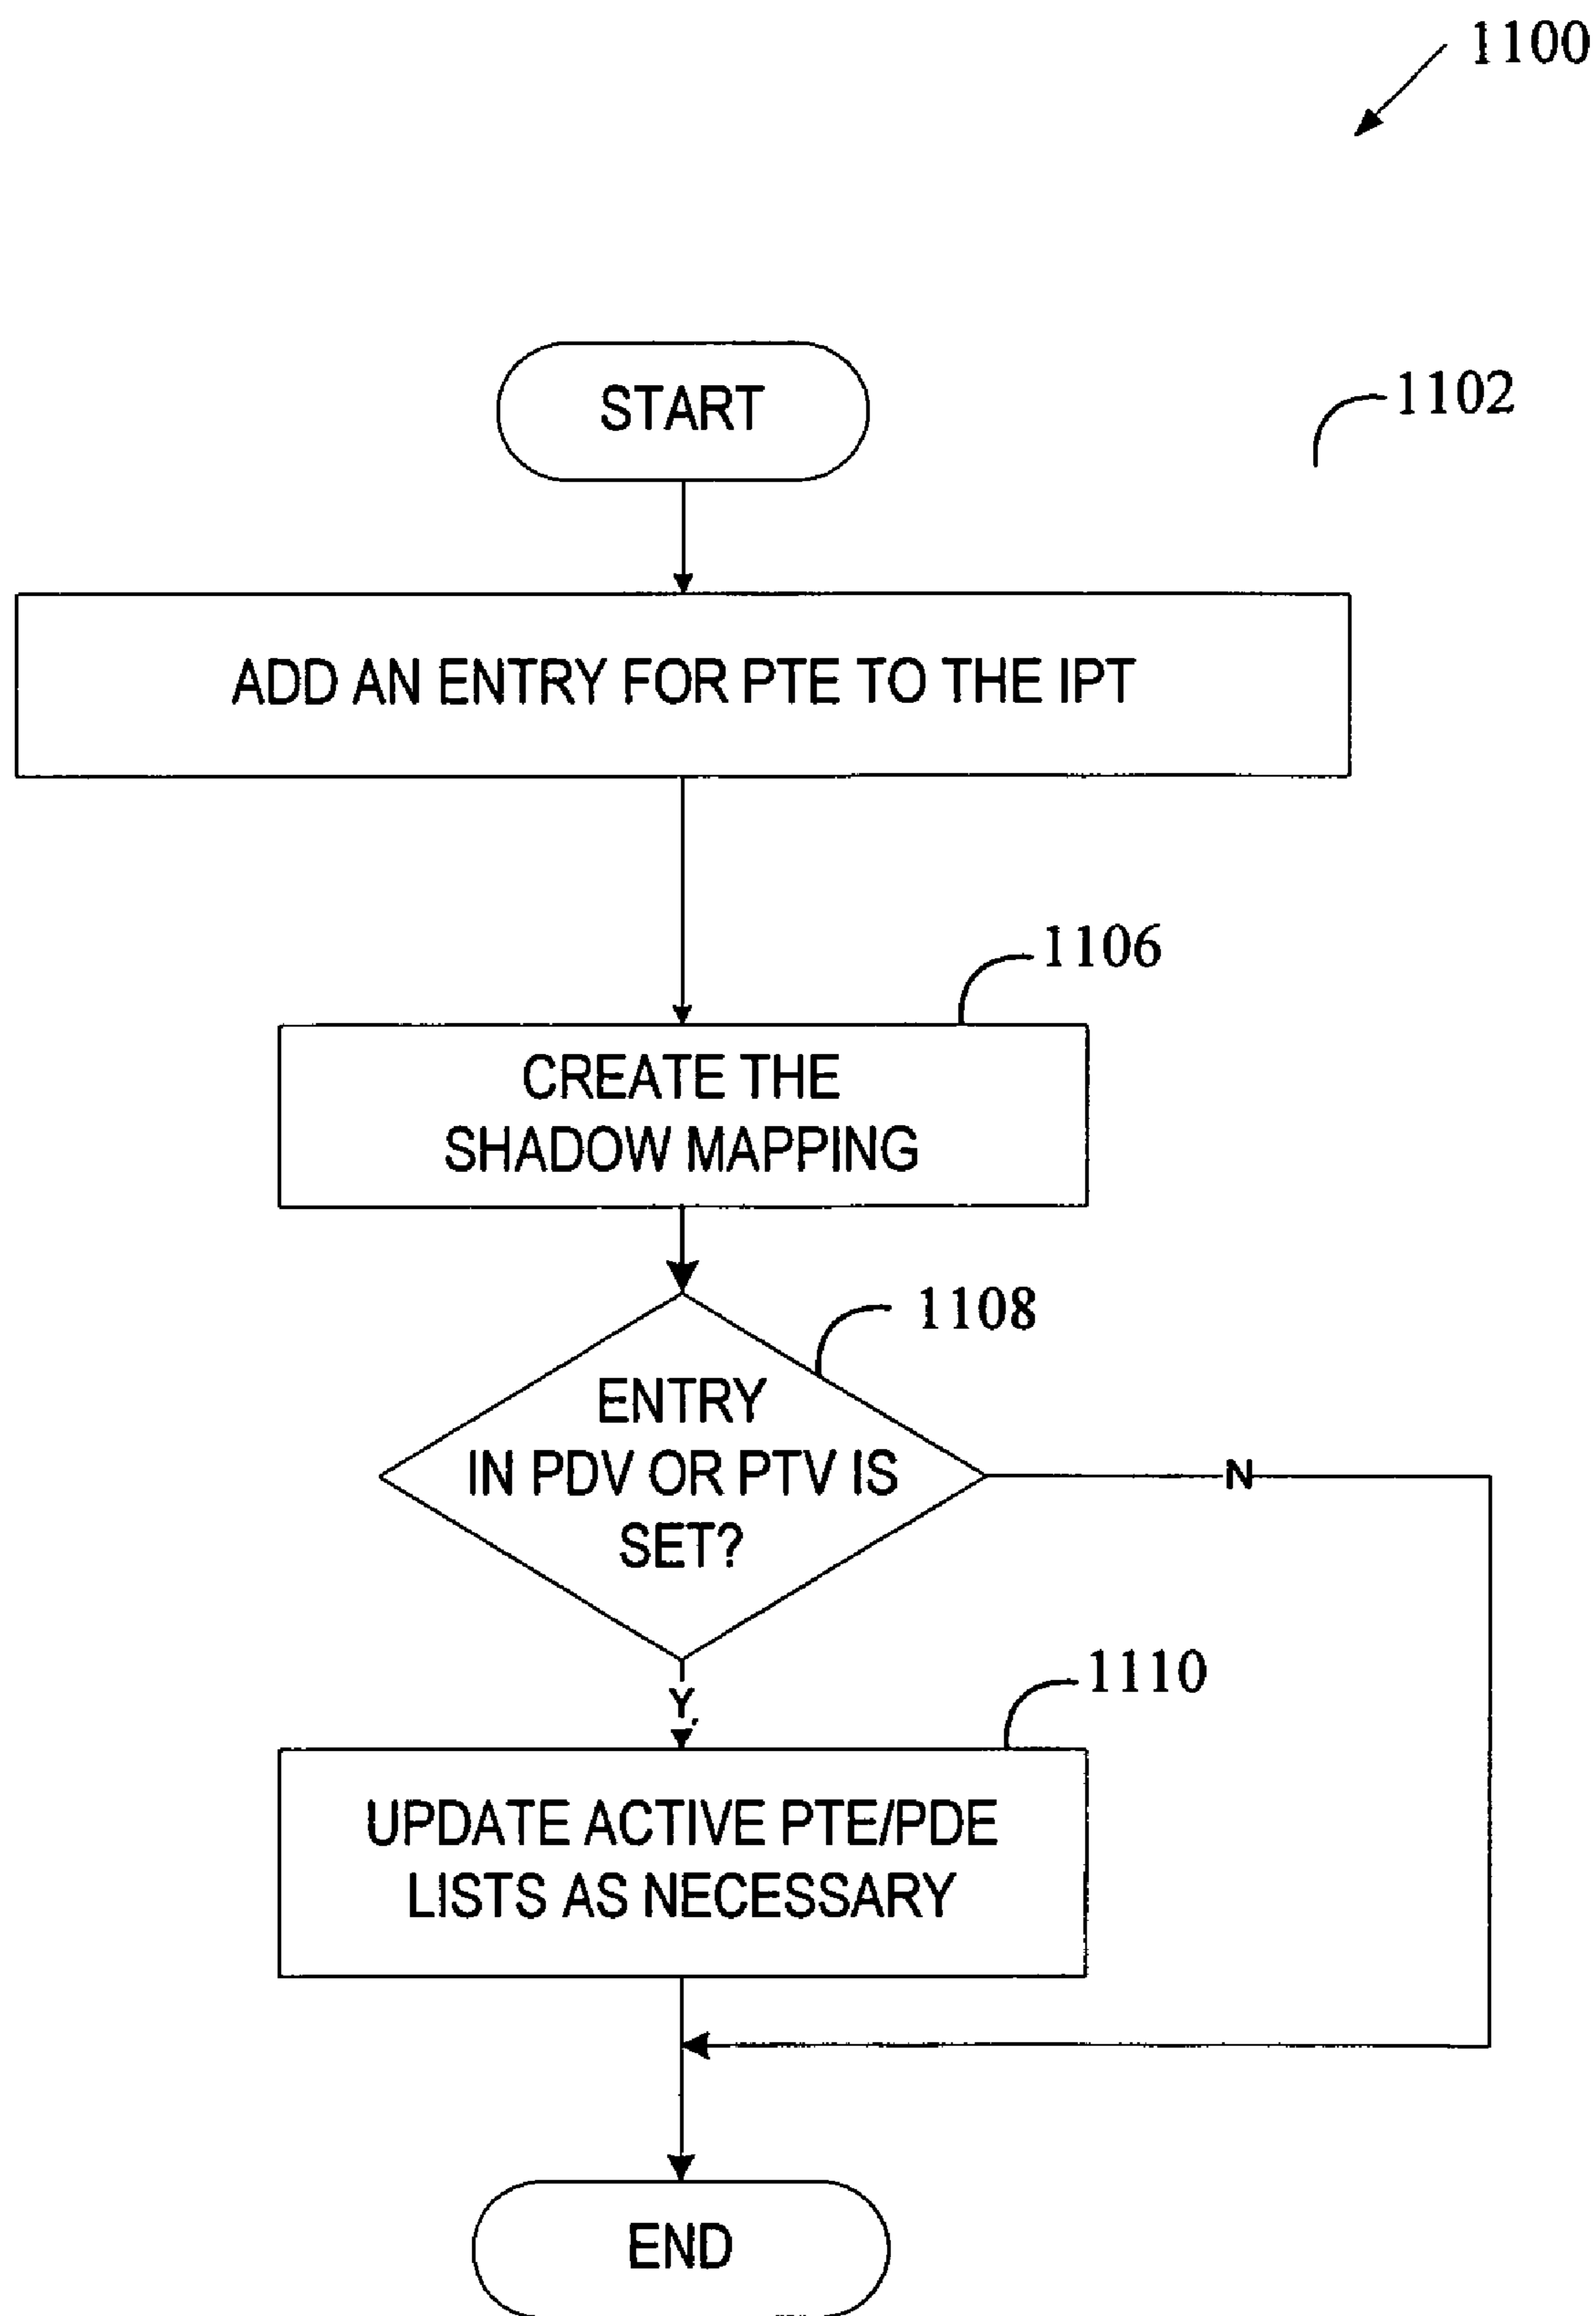


FIG. 11

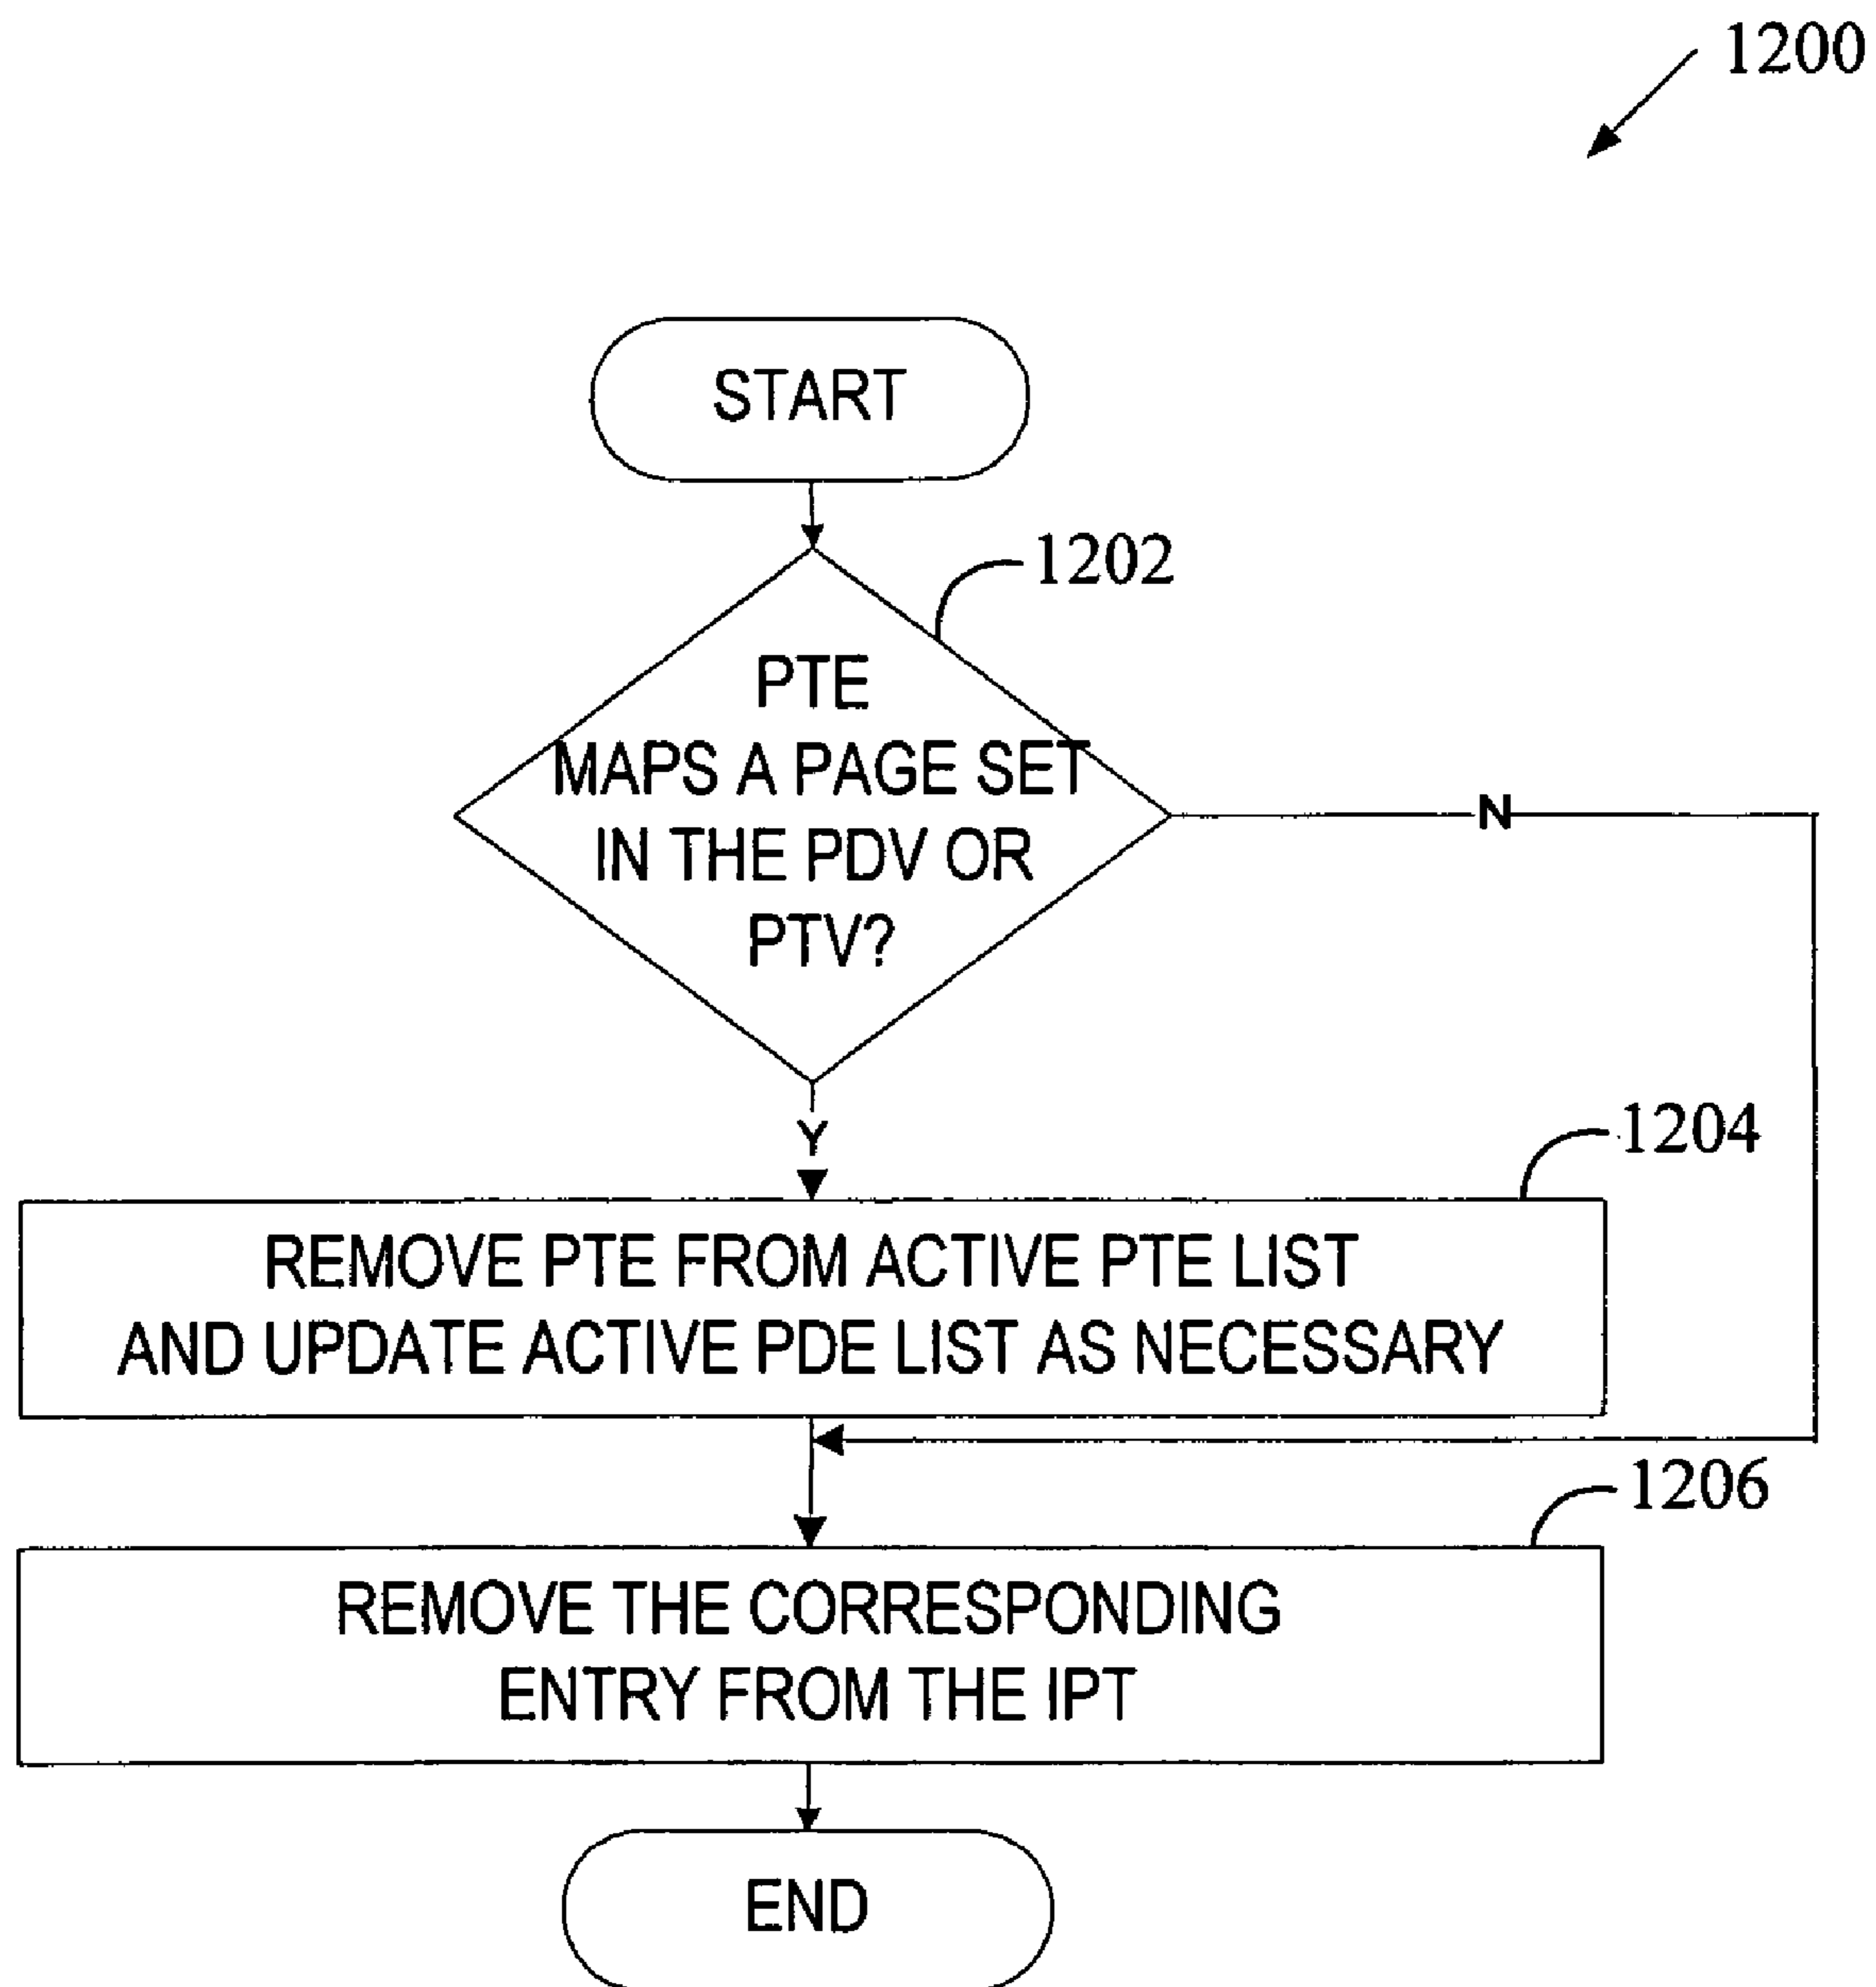


FIG. 12

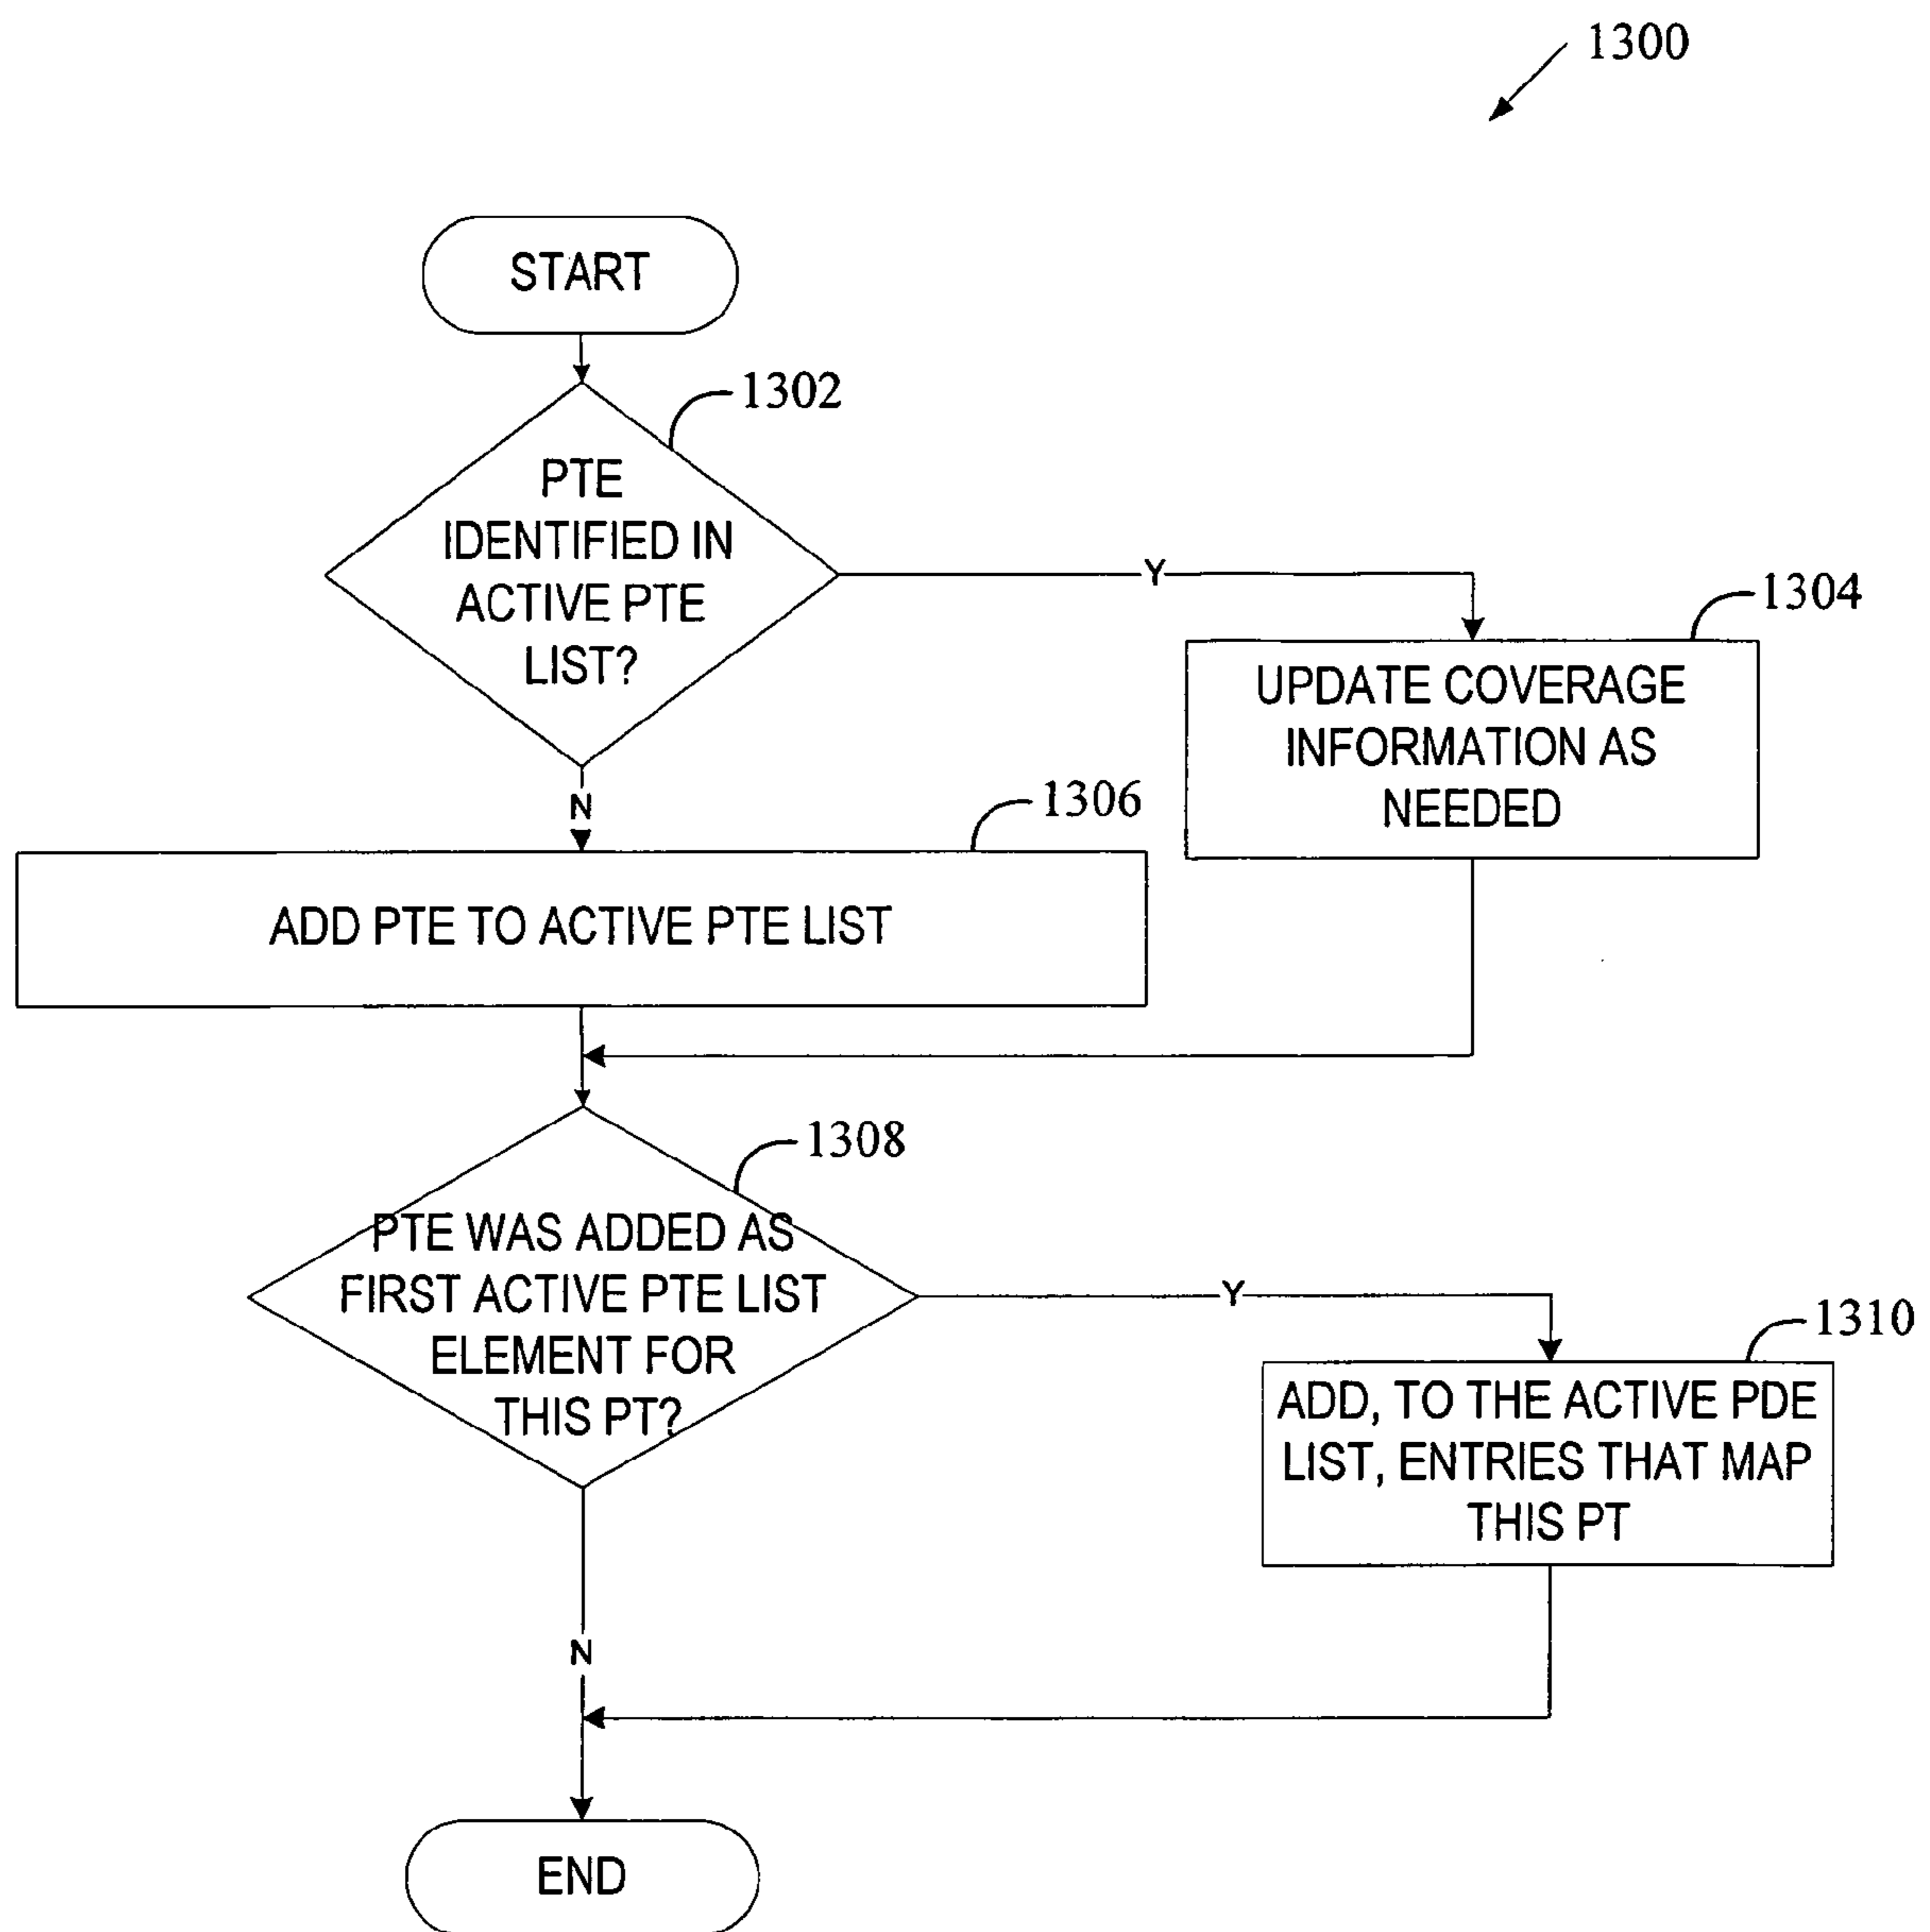


FIG. 13




FIG. 14

1

METHOD AND APPARATUS FOR SUPPORTING ADDRESS TRANSLATION IN A VIRTUAL MACHINE ENVIRONMENT

FIELD

Embodiments of the invention relate generally to virtual machines, and more specifically to supporting address translation in a virtual machine environment.

BACKGROUND

A conventional virtual-machine monitor (VMM) typically runs on a computer and presents to other software the abstraction of one or more virtual machines. Each virtual machine may function as a self-contained platform, running its own “guest operating system” (i.e., an operating system (OS) hosted by the VMM) and other software, collectively referred to as guest software. The guest software expects to operate as if it were running on a dedicated computer rather than a virtual machine. That is, the guest software expects to control various events and have access to hardware resources such as physical memory and memory-mapped input/output (I/O) devices. For example, the guest software expects to maintain control over address-translation operations and have the ability to allocate physical memory, provide protection from and between guest applications, use a variety of paging techniques, etc. However, in a virtual-machine environment, the VMM should be able to have ultimate control over the computer’s resources to provide protection from and between virtual machines.

BRIEF DESCRIPTION OF THE DRAWINGS

The present invention is illustrated by way of example, and not by way of limitation, in the figures of the accompanying drawings and in which like reference numerals refer to similar elements and in which:

FIG. 1 illustrates one embodiment of a virtual-machine environment, in which the present invention may operate;

FIG. 2 illustrates operation of a virtual TLB, according to one embodiment of the present invention;

FIGS. 3A and 3B illustrate a process of creating and maintaining metadata for a shadow PT hierarchy, according to two alternative embodiments of the present invention;

FIG. 4 is a flow diagram of one embodiment of a process for synchronizing guest translation data structure and shadow translation data structure;

FIG. 5 is a flow diagram of one embodiment of a process for maintaining metadata for a shadow translation data structure;

FIG. 6 is a flow diagram of one embodiment of a process for facilitating a change of an address space;

FIG. 7 is a flow diagram of one embodiment of a process for synchronizing entries of two translation data structures for a specified address;

FIG. 8 is a flow diagram of one embodiment of a process for removing a shadow PT hierarchy from a working set of shadow PT hierarchies maintained by the VMM;

FIG. 9 is a flow diagram of one embodiment of a process for adding an entry to a PD of a shadow PT hierarchy;

FIG. 10 is a flow diagram of one embodiment of a process for removing an entry from a PD of a shadow PT hierarchy;

FIG. 11 is a flow diagram of one embodiment of a process for adding an entry to a PT of a shadow PT hierarchy;

FIG. 12 is a flow diagram of one embodiment of a process for removing an entry from a PT of a shadow PT hierarchy;

2

FIG. 13 is a flow diagram of one embodiment of a process for monitoring a PTE of a shadow PT hierarchy; and

FIG. 14 is a flow diagram of one embodiment of a process for removing monitoring from a PTE of a shadow PT hierarchy.

DESCRIPTION OF EMBODIMENTS

A method and apparatus for supporting address translation in a virtual machine environment is described. In the following description, for purposes of explanation, numerous specific details are set forth in order to provide a thorough understanding of the present invention. It will be apparent, however, to one skilled in the art that the present invention can be practiced without these specific details.

Some portions of the detailed descriptions that follow are presented in terms of algorithms and symbolic representations of operations on data bits within a computer system’s registers or memory. These algorithmic descriptions and representations are the means used by those skilled in the data processing arts to convey most effectively the substance of their work to others skilled in the art. An algorithm is here, and generally, conceived to be a self-consistent sequence of operations leading to a desired result. The operations are those requiring physical manipulations of physical quantities. Usually, though not necessarily, these quantities take the form of electrical or magnetic signals capable of being stored, transferred, combined, compared, and otherwise manipulated. It has proven convenient at times, principally for reasons of common usage, to refer to these signals as bits, values, elements, symbols, characters, terms, numbers, or the like.

It should be borne in mind, however, that all of these and similar terms are to be associated with the appropriate physical quantities and are merely convenient labels applied to these quantities. Unless specifically stated otherwise as apparent from the following discussions, it is appreciated that throughout the present invention, discussions utilizing terms such as “processing” or “computing” or “calculating” or “determining” or the like, may refer to the action and processes of a computer system, or similar electronic computing device, that manipulates and transforms data represented as physical (electronic) quantities within the computer system’s registers and memories into other data similarly represented as physical quantities within the computer-system memories or registers or other such information storage, transmission or display devices.

In the following detailed description of the embodiments, reference is made to the accompanying drawings that show, by way of illustration, specific embodiments in which the invention may be practiced. In the drawings, like numerals describe substantially similar components throughout the several views. These embodiments are described in sufficient detail to enable those skilled in the art to practice the invention. Other embodiments may be utilized and structural, logical, and electrical changes may be made without departing from the scope of the present invention. Moreover, it is to be understood that the various embodiments of the invention, although different, are not necessarily mutually exclusive. For example, a particular feature, structure, or characteristic described in one embodiment may be included within other embodiments. The following detailed description is, therefore, not to be taken in a limiting sense, and the scope of the present invention is defined only by the appended claims, along with the full scope of equivalents to which such claims are entitled.

Although the below examples may describe providing support for address translation in a virtual machine environment

in the context of execution units and logic circuits, other embodiments of the present invention can be accomplished by way of software. For example, in some embodiments, the present invention may be provided as a computer program product or software which may include a machine or computer-readable medium having stored thereon instructions which may be used to program a computer (or other electronic devices) to perform a process according to the present invention. In other embodiments, processes of the present invention might be performed by specific hardware components that contain hardwired logic for performing the processes, or by any combination of programmed computer components and custom hardware components.

Thus, a machine-readable medium may include any mechanism for storing or transmitting information in a form readable by a machine (e.g., a computer), but is not limited to computer-readable storage media such as floppy diskettes, optical disks, Compact Disc, Read-Only Memories (CD-ROMs), and magneto-optical disks, Read-Only Memories (ROMs), Random Access Memories (RAMs), Erasable Programmable Read-Only Memories (EPROMs), Electrically Erasable Programmable Read-Only Memories (EEPROMs), magnetic or optical cards, flash memories, and transmission media such as a transmission over the Internet, electrical, optical, acoustical or other forms of propagated signals (e.g., carrier waves, infrared signals, digital signals, etc.) or the like.

Further, a design may go through various stages, from creation to simulation to fabrication. Data representing a design may represent the design in a number of manners. First, as is useful in simulations, the hardware may be represented using a hardware description language or another functional description language. Additionally, a circuit level model with logic and/or transistor gates may be produced at some stages of the design process. Furthermore, most designs, at some stage, reach a level of data representing the physical placement of various devices in the hardware model. In the case where conventional semiconductor fabrication techniques are used, data representing a hardware model may be the data specifying the presence or absence of various features on different mask layers for masks used to produce the integrated circuit. In any representation of the design, the data may be stored in any form of a machine-readable medium. An optical or electrical wave modulated or otherwise generated to transmit such information, a memory, or a magnetic or optical storage such as a disc may be the machine readable medium. Any of these mediums may “carry” or “indicate” the design or software information. When an electrical carrier wave indicating or carrying the code or design is transmitted, to the extent that copying, buffering, or re-transmission of the electrical signal is performed, a new copy is made. Thus, a communication provider or a network provider may make copies of an article (a carrier wave) embodying techniques of the present invention.

FIG. 1 illustrates one embodiment of a virtual-machine environment **100**, in which the present invention may operate. In this embodiment, bare platform hardware **116** comprises a computing platform, which may be capable, for example, of executing a standard operating system (OS) or a virtual-machine monitor (VMM), such as a VMM **112**.

The VMM **112**, typically implemented in software, may emulate and export a bare machine interface to higher level software. Such higher level software may comprise a standard or real-time OS, may be a highly stripped-down operating environment with limited operating system functionality, may not include traditional OS facilities, etc. Alternatively, for example, the VMM **112** may be run within, or on top of,

another VMM. VMMs may be implemented, for example, in hardware, software, firmware or by a combination of various techniques.

The platform hardware **116** can be of a personal computer (PC), mainframe, handheld device, portable computer, set-top box, or any other computing system. The platform hardware **116** includes a processor **118** and memory **120**.

Processor **118** can be any type of processor capable of executing software, such as a microprocessor, digital signal processor, microcontroller, or the like. The processor **118** may include microcode, programmable logic or hard-coded logic for performing the execution of method embodiments of the present invention. Although FIG. 1 shows only one such processor **118**, there may be one or more processors in the system.

Memory **120** can be a hard disk, a floppy disk, random access memory (RAM) (e.g., dynamic RAM (DRAM) or static RAM (SRAM)), read only memory (ROM), flash memory, any combination of the above devices, or any other type of machine medium readable by processor **118**. Memory **120** may store instructions and/or data for performing the execution of method embodiments of the present invention.

The VMM **112** presents to other software (i.e., “guest” software) the abstraction of one or more virtual machines (VMs), which may provide the same or different abstractions to the various guests. FIG. 1 shows two VMs, **102** and **114**. The guest software running on each VM may include a guest OS such as a guest OS **104** or **106** and various guest software applications **108** and **110**. Each of the guest OSs **104** and **106** expects to access physical resources (e.g., processor registers, memory and I/O devices) within the VMs **102** and **114** on which the guest OS **104** or **106** is running and to perform other functions. For example, the guest OS **104** or **106** expects to have access to all registers, caches, structures, I/O devices, memory and the like, according to the architecture of the processor and platform presented in the VM **102** and **114**. The resources that can be accessed by the guest software may either be classified as “privileged” or “non-privileged.” For privileged resources, the VMM **112** facilitates functionality desired by guest software while retaining ultimate control over these privileged resources. Non-privileged resources do not need to be controlled by the VMM **112** and can be accessed directly by guest software.

Further, each guest OS expects to handle various fault events such as exceptions (e.g., page faults, general protection faults, etc.), interrupts (e.g., hardware interrupts, software interrupts), and platform events (e.g., initialization (INIT) and system management interrupts (SMIs)). Some of these fault events are “privileged” because they must be handled by the VMM **112** to ensure proper operation of VMs **102** and **114** and for protection from and among guest software.

When a privileged fault event occurs or guest software attempts to access a privileged resource, control may be transferred to the VMM **112**. The transfer of control from guest software to the VMM **112** is referred to herein as a VM exit. After facilitating the resource access or handling the event appropriately, the VMM **112** may return control to guest software. The transfer of control from the VMM **112** to guest software is referred to as a VM entry.

In one embodiment, the processor **118** controls the operation of the VMs **102** and **114** in accordance with data stored in a virtual machine control structure (VMCS) **125**. The VMCS **125** is a structure that may contain state of guest software, state of the VMM **112**, execution control information indicating how the VMM **112** wishes to control operation of guest software, information controlling transitions between the

VMM 112 and a VM, etc. The processor 118 reads information from the VMCS 125 to determine the execution environment of the VM and to constrain its behavior. In one embodiment, the VMCS is stored in memory 120. In some embodiments, multiple VMCS structures are used to support multiple VMs.

During address translation operations, the VM 102 or 114 expects to allocate physical memory, provide protection from and between guest software applications (e.g., applications 108 or 110), use a variety of paging techniques, etc. In a non-virtual machine environment, an address translation mechanism expected by an OS may be based on a translation lookaside buffer (TLB) 122 controlled by the processor 118 and a translation data structure, such as a page-table (PT) hierarchy, controlled by the OS and used to translate virtual memory addresses into physical memory addresses when paging is enabled.

The architecture of the Intel® Pentium® 4 Processor supports a number of paging modes. The most commonly used paging mode supports a 32-bit linear address space using a two-level hierarchical paging structure (referred to herein as a two-level hierarchy paging mode). Embodiments of the invention are not limited to this paging mode, but instead may be employed by one skilled in the art to virtualize other paging modes (e.g., Physical Address Extension (PAE) mode, Intel® Extended Memory 64 Technology (EM64T) mode, etc.) and implementations (e.g., hashed page tables). In one embodiment based on a TLB, translation of a virtual memory address into a physical memory address begins with searching the TLB 122 using either the upper 20 bits (for a 4 KB page frame) or the upper 10 bits (for a 4 MB page frame) of the virtual address. If a match is found (a TLB hit), the upper bits of a physical page frame that are contained in the TLB 122 are conjoined with the lower bits of the virtual address to form a physical address. The TLB also contains access and permission attributes associated with the mapping. If no match is found (a TLB miss), the processor consults the PT hierarchy to determine the virtual-to-physical translation, which is then cached in the TLB 122. Entries in the PT hierarchy may include some attributes that are automatically set by the processor on certain accesses.

If the PT hierarchy is modified, the TLB 122 may become inconsistent with the PT hierarchy if a corresponding address translation exists in the TLB 122. The OS may expect to be able to resolve such an inconsistency by issuing an instruction to the processor 118. For example, in the instruction set architecture (ISA) of the Intel® Pentium® 4 (referred to herein as the IA-32 ISA), a processor allows software to invalidate cached translations in the TLB by issuing the INVLPG instruction. In addition, the OS may expect to request the processor 118 to change the address space completely, which should result in the removal of all translations from the TLB 122. For example, in the IA-32 ISA, an OS may use a MOV instruction or a task switch to request a processor to load CR3 (which contains the base address of the PT hierarchy), thereby removing all translations from the TLB. Different levels of the page table hierarchy may have different names based upon mode and implementation. In the two-level hierarchy paging mode, there are two levels of paging structures. The CR3 register points to the base of the page directory page. Entries in the page directory may either specify a mapping to a large-size page (e.g., a 4 MB superpage, a 2 MB superpage, 1 GB superpage, etc.), or a reference to a page table. The page table in turn may contain mappings to small-size pages.

As discussed above, in the virtual-machine environment, the VMM 112 should be able to have ultimate control over physical resources including the TLB 122. Embodiments of

the present invention address the conflict between the expectations of the VMs 102 and 114 and the role of the VMM 112 by using a virtual TLB that emulates the functionality of the processor's physical TLB.

The virtual TLB includes the TLB 122 and a set of shadow PT hierarchies controlled by the VMM 112. The set of shadow PT hierarchies derive its format and content from guest PT hierarchies that may be currently used or not used by the VM 102 or 114. If the VM 102 or 114 modifies the content of the guest PT hierarchies, this content becomes inconsistent with the content of the shadow PT hierarchies. The inconsistencies between the guest PT hierarchies and the shadow PT hierarchies are resolved using techniques analogous to those employed by the processor 118 in managing the TLB 122. Some of these techniques force the VM 102 or 114 to issue an event indicating an attempt to manipulate the TLB (e.g., INVLPG, page fault, and load CR3). Such events are privileged and, therefore, result in a VM exit to the VMM 112. The VMM then evaluates the event and synchronizes all maintained shadow PT hierarchies with the current guest state if needed. We will refer to the set of maintained shadow PT hierarchies as the working set. As multiple processes may use the same guest page table, it is possible for the same shadow PT to be a part of multiple guest PT hierarchies. The corresponding shadow PT will in turn be a member of multiple shadow PT hierarchies.

Note that synchronization performed by the VMM may update shadow page table or page directory entries for a shadow PT hierarchy that is not currently in-use. Likewise synchronization may be required to guest pages that are not part of the in-use guest PT hierarchy.

In one embodiment, the VMM 112 includes an address translation module 126 that is responsible for creating and maintaining a working set of shadow PT hierarchies for each of the VM 102 and 114 in a virtual TLB (VTLB) data store 124. The working set of shadow PT hierarchies is maintained for corresponding active processes of the VM 102 or 114 (i.e., processes that are likely to be activated in the near future by the VM 102 or 114). With the IA32 ISA, the only explicitly defined guest hierarchy is that defined by the currently used paging structures. In practice there is a high deal of temporal locality for guest processes and their address spaces. The VMM may employ heuristics or explicit information to determine a set of active process.

When the VM 102 or 114 enables a guest PT hierarchy for one of the active processes of the VM 102 or 114, the address translation module 126 identifies a corresponding shadow PT hierarchy in the working set and requests the processor 118 to load its base address. When applicable, the address translation module 126 can then reuse previously computed mappings that are stored in the shadow PT hierarchies.

If the VM 102 or 114 activates a new process, the address translation module 126 derives a new shadow PT hierarchy from a corresponding guest PT hierarchy and adds it to the working set. Alternatively, if the VM 102 or 114 de-activates an existing process, the address translation module 126 removes information corresponding to the guest PT hierarchy from the working set.

In one embodiment, the address translation module 126 is responsible for extracting metadata from each new shadow PT hierarchy, storing the metadata in the VTLB data store 124, and updating the metadata when the shadow PT hierarchy is modified. In one embodiment, the metadata includes a PT vector (PTV), a PD vector (PDV), an active PTE list, and an active PDE list.

The PTV and PDV track the guest frames that are used as PTs and PDs. In one embodiment, this information is encoded

in bit vectors. The PTV may be indexed by page frame number (PFN), with each entry bit being set if a corresponding PFN is a PT. The PDV may be indexed by a page frame number (PFN), with each entry bit being set if a corresponding PFN is a PD.

The active PTE list is a list of PT entries (PTEs) in the shadow PT hierarchy that point to frames holding PTs and PD. The active PDE list identifies PD entries (PDEs) in the shadow PT hierarchy that point to PTs containing PT entries identified in the active PTE list.

In one embodiment, active PDE and PTE lists contain additional metadata describing whether the mapping is to a PD or PT frame.

One skilled in the art will understand that embodiments of this invention may use a variety of data structures which may be more or less space or time efficient than those described herein. One skilled in the art will also recognize the extension of tracking structures to support additional paging modes. For example, an EM64T paging mode maps a 64-bit virtual address to a physical address through a four-level hierarchical paging structure. The actual number of bits supported in the virtual or physical address spaces may be implementation dependent and may be less than 64 bits in a particular implementation. As will be discussed in more detail below, an EM64T implementation may require additions of a page-map level 4 (PML4) page vector and a page directory pointer (PDP) page vector to track the additional page tables used in the EM64T paging structure. Likewise, one skilled in the art will recognize that the active PTE list will be extended to include entries which map any page used within the paging structures (e.g., PML4 or PDP pages for EM64T).

In one embodiment, active PTE/PDE list metadata is maintained to track the number of PD and PT frames that are mapped through a page table. When the number of mappings per page is incremented from 0, then PDEs which map the PT must be added to the active PDE list, and when the number of mappings is decreased to zero, then PDEs that map this PT must be removed from the active PDE list.

In one embodiment, the address translation module 126 is responsible for synchronizing a current shadow PT hierarchy with a current guest PT hierarchy when such synchronization is needed. The address translation module 126 performs the synchronization by determining which entries in the guest PT hierarchy have recently been modified and then updating corresponding entries in the shadow PT hierarchy accordingly. The address translation module 126 determines which entries in the guest PT hierarchy have recently been modified based on the metadata extracted from the shadow PT hierarchy and attributes associated with the entries of the shadow PT hierarchy. In one embodiment, the attributes include access attributes associated with PD entries in the shadow PT hierarchy and update attributes associated with PT entries in the shadow PT hierarchy.

FIG. 2 illustrates operation of a virtual TLB 204, according to one embodiment of the present invention. Virtual TLB 204 includes a shadow translation data structure represented by a shadow PT hierarchy 206 and a physical TLB 208. The shadow PT hierarchy 206 derives its structure and content from a guest translation data structure represented by a guest PT hierarchy 202. In one embodiment, the VMM maintains a working set of shadow PT hierarchies for active processes of the VM.

In one embodiment, when the VM requests the processor to enable a different guest PT hierarchy (e.g., by issuing MOV to CR3 or task switch in the IA-32 ISA), control transitions to the VMM, which instructs the processor to load the base address 214 of a shadow PT hierarchy 206 corresponding to

the requested guest PT hierarchy 202. In some embodiments, this shadow PT hierarchy 206 is synchronized with the guest PT hierarchy 202 using relevant metadata and attributes, as will be discussed in greater detail below.

In one embodiment, the virtual TLB maintains access and update attributes in the entries of the shadow PD and PTs. These attributes are also referred to as an accessed (A) bit and a dirty (D) bit. In one embodiment, when a page frame is accessed by guest software for the first time, the processor sets the accessed (A) attribute in the corresponding PT entry or PD entry in the shadow PT hierarchy 206. If guest software attempts to write a page frame, the processor sets the dirty (D) attribute in the corresponding shadow PT entry.

Guest software is allowed to freely modify the guest PT hierarchy 202 including changing virtual-to-physical mapping, permissions, etc. Accordingly, the shadow PT hierarchy 206 may not be always consistent with the guest PT hierarchy 202. When a problem arises from an inconsistency between the hierarchies 202 and 206, the guest OS, which treats the virtual TLB 204 as a physical TLB, attempts to change the virtual TLB 204 by requesting a processor to perform an operation defined by a relevant ISA. For example, in the IA-32 ISA, such operations include the INVLPG instruction, CR3 loads, paging activation (modification of CR0.PG), modification of global paging (toggling of the CR4.PGE bit), etc. The operations attempting to change the virtual TLB 204 are configured by the VMM as privileged (e.g., using corresponding execution controls stored in the VMCS), and, therefore, result in a VM exit to the VMM. The VMM then determines the cause of the VM exit and modifies the content of the shadow PT hierarchy 206 if necessary. For example, if the VM exit occurs due to a page fault that should be handled by the guest OS (e.g., a page fault caused by an access not permitted by the guest PT hierarchy 202), the page fault is injected to the guest OS for handling. Alternatively, if the VM exit occurs due to a page fault (or any other operations such as INVLPG) resulting from an inconsistency between the entries of the hierarchies 202 and 206, the VMM may need to remove stale entries, add new entries, or modify existing entries, as will be discussed in more detail below. Page faults caused by the guest PT hierarchy are referred to herein as 'real' page faults, and page faults that would not have occurred with direct usage of the guest page tables are referred to herein as 'induced' page faults.

FIG. 3A illustrates a process of creating and maintaining metadata for a shadow PT hierarchy in a two-level hierarchy paging mode, according to one embodiment of the present invention.

Referring to FIG. 3A, a number of physical page frames identified by distinct letters (letters A through W) is illustrated. Some guest page frames may contain a PD (e.g., frame A). Other guest page frames may contain a PT (e.g., frames B, C, and L). A hierarchy 302 is a guest PT hierarchy.

FIG. 3A shows a shadow PT hierarchy 304 created based on a guest PT hierarchy 302. Each PD or PT in the guest PT hierarchy 302 includes a corresponding PD or PT in the shadow PT hierarchy 304. Note that in general a shadow page is not required for each page in the guest PT. Some embodiments may choose to restrict shadow pages according to usage statistics (e.g., only generate shadow pages for guest PT pages that have been used), or according to resource constraints (e.g., maintaining only a set of shadow pages based on available memory). Separate shadow tables are maintained for PD and PT tables derived from the same physical frame. For example, separate tables 330 and 332 are maintained for PD 306 and PT 308 that are derived from the same physical frame 314. The PD and PT entries in the

shadow PT hierarchy **304** contain transformed mappings for the guest frames **314** through **324**.

In the guest PT hierarchy **302**, frames **316** and **318** are used as PTs **310** and **312**, and frame **314** is used both as PD **306** and PT **308**. This usage is illustrated as “PT” and “PD/PT” in the page frames **314** through **316** shown under the shadow PT hierarchy **304**.

The shadow PT hierarchy **304** is associated with an active PTE list **342** and an active PDE list **344**. In one embodiment, the active PTE list **342** identifies PT entries in the shadow PT hierarchy **304** that map PT and PD page frames from the guest PT hierarchy **302**. In particular, the active PTE list **342** identifies entries in the PT **332** that map page frames **314** through **318**. In one embodiment, the active PDE list **344** identifies PD entries in the shadow PT hierarchy that point to PTs with entries identified in the active PTE list **342**. In particular, the active PDE list **344** includes entries in the PD **330** that point to the PT **332**. The active PTE list **342** and the active PDE list **344** are components of the metadata of the shadow PT hierarchy **304**.

The shadow PT hierarchy **304** is associated with a PT bit vector (PTV) **362** and a PD bit vector (PDV) **364**. In one embodiment, the PTV **362** tracks the guest page frames that are used as PTs. In particular, the PTV **362** includes page frames **314** through **318** which are used as PTs in the guest PT hierarchy **302**. In one embodiment, the PDV **364** tracks the guest page frames that are used as PDs. In particular, the PDV **364** includes page frame **314** that is used as PD in the guest PT hierarchy **302**. In one embodiment, the PTV **362** and PDV **364** represent all shadow PT hierarchies in the working set and track the capacity in which shadow pages are employed in the working set (e.g., if a shadow page has not been allocated for a guest PT, then the PTV will not reflect the guest PT page, even if it appears in the guest paging structures).

In one embodiment, if the guest OS adds a new PT to the guest PT hierarchy **302**, the VMM may detect this addition (e.g., on the next or subsequent VM exit related to TLB manipulation) and add a corresponding PT to the shadow PT hierarchy **304**. For example, if a new PT **352** derived from a frame **319** is added to the guest PT hierarchy **302**, with a mapping for a new frame **354**, the VMM may add a corresponding PT **360** with transformed mappings to the shadow PT hierarchy **304** and update the metadata to reflect this change. In particular, the VMM adds an entry mapping frame **319** in the PT **332** to the active PTE list **342**, and an entry pointing to the PT **360** in the PD **330** to the active PDE list **344**. Also, the VMM adds frame **319** to PTV **362**, which tracks guest frames (i.e., here frame **319**) used as PTs.

FIG. **3B** illustrates a process of creating and maintaining metadata for a shadow PT hierarchy in the EM64T paging mode, according to one embodiment of the present invention.

Referring to FIG. **3B**, the base of the paging structure is a PML4 page (e.g., frame A). Each entry in the PML4 page may reference a PDP page (e.g., frames B and C). Each entry in the PDP page may reference a page directory (PD) page (e.g., frame D or E), each entry of which in turn may reference a page in a page table (PT) page (e.g., frame F, G, H or I).

Each PML4, PDP, PD or PT page may be 4 KB in size. In order to support physical address spaces larger than 32 bits, the entry size may be increased relative to the 32-bit paging mode. Specifically, there may be 512 entries per page, requiring that 9 bits of the virtual address be used at each level to select the appropriate entry. This selector size may lead to a large page size of 2 MB instead of 4 MB as described previously.

In one embodiment, the creation of metadata in the EM64T paging mode includes the generation of several vectors, an

active entry list, and several active directory lists. The vectors include a PML4V vector identifying frames used as PML4 pages, a PDPV vector identifying frames used as PDP pages, a PDV vector identifying frames used as PD pages, and a PTV vector identifying frames used as PT pages. The active entry list is an active PTE list including all mappings which map a PML4, PDP, PD or PT page. The active directory lists include lists identifying higher level mapping structures referencing a lower level structure through which the guest page corresponding to a shadow structure can be accessed. In particular, the active directory lists consist of an active PDE list including those PDEs that reference a page containing active PTE list entries, an active PDPE list including active PDPE entries which reference a PD containing an active PDE list entry, and an active PML4E list including entries which map a PDP containing elements in the active PDPE list.

In one embodiment, the synchronization of the shadow page tables begins with checking each entry in the active PML4E list associated with the used shadow PT hierarchy. If the entry has been accessed, each element in the active PDPE list corresponding to the accessed PML4 entry is checked, and then the processing continues as previously described.

In an alternative embodiment, active lists are not maintained and/or processed for one or more of the upper levels of the hierarchy. For example, in a system in which only a single entry is populated in the uppermost paging structure, the use of an active list for each level of the hierarchy will cause this single entry to be always accessed, thereby allowing no reduction in the amount of processing required for lower levels in the hierarchy. To accommodate this usage model, the synchronization may instead begin by processing an active list lower in the hierarchy. For example, in one embodiment, active PDPE list elements may first be processed followed by active PDE list elements or active PTE list elements associated with a used shadow PT hierarchy. In one embodiment, the initial layer processed on synchronization may be predetermined. In another embodiment, the initial layer to be processed may be determined by dynamic profiling of the guest's page table usage.

Various other paging modes may be used with embodiments of the present invention. For example, IA-32 supports an additional paging mode in which a 32-bit virtual address is mapped to a larger physical address. In this additional mode of operation, the page table base register is configured to point to a PDP page which contains four elements. Entry sizes and behaviors in this additional mode of operations are similar to those described above for the 64-bit virtual address mode. As this additional mode does not make use of PML4 pages, the PML4V and active PML4E list are not required.

FIG. **4** is a flow diagram of one embodiment of a process **400** for synchronizing a guest translation data structure and a shadow translation data structure. The process may be performed by processing logic that may comprise hardware (e.g., dedicated logic, programmable logic, microcode, etc.), software (such as that run on a general purpose computer system or a dedicated machine), or a combination of both. In one embodiment, the process is performed by an address translation module **126** of FIG. **1**.

Referring to FIG. **4**, process **400** begins with processing logic receiving control transitioned from a VM due to an event pertaining to manipulation of the TLB (processing block **402**). Examples of such events may include a request to change the current address space (e.g., CR3 load), a request to adjust inconsistent translations for a specified virtual address in the TLB (e.g., INVLPG), a page fault, etc.

At processing box **404**, processing logic determines whether the event pertaining to the manipulation of the TLB

should be handled by the VM. If so (e.g., the event is a page fault caused by a problematic mapping in a guest translation data structure), control is returned to the VM for handling the event (processing block **406**). If not, processing logic determines whether the event is associated with a specified problematic address (processing box **408**).

If the event does not need to be handled by the VM, it may be associated with a specified problematic address. Examples of such an event may include an event caused by the INVLPG instruction that takes a problematic address as an operand, an event caused by an induced page fault (e.g., a page fault resulting from an inconsistency between the two translation data structures with respect to a specific mapping, a page fault caused by a need to virtualize A/D bits in the guest translation data structure, etc.), etc. If the event is associated with a specified problematic address, processing logic makes corrections in the shadow translation data structure for the specified address (e.g., removes a stale mapping for the specified address or adds a new mapping for the specified address) to conform to the guest translation data structure (processing block **410**). One embodiment of a process for synchronizing entries of two translation data structures for a specified address is discussed in more detail below in conjunction with FIG. 7.

If the event is not associated with any specific address (e.g., the event is caused by a request of the VM to change the address space, which flushes all TLB entries in IA32), processing logic determines which entries of the guest translation data structure have been modified (processing block **412**). The determination is made using metadata extracted from the shadow translation data structure and attributes associated with the entries of the shadow translation data structure (processing block **412**). The metadata includes vectors and active lists for various levels of the shadow translation data structure. A vector for a specific level of the shadow translation data structure identifies frames used as pages at this level of the guest translation data structure. The active lists include an active entry list and one or more active directory lists. The active entry list includes mappings that map pages used by the guest in forming the guest translation data structure. The active directory lists identify higher level mapping structures referencing a lower level structure through which a guest page corresponding to a shadowed paging structure can be accessed. As discussed above, in the two-level hierarchy paging mode, the metadata includes, in one embodiment, vectors PTV and PDV, an active entry list (a PTE list), and an active directory list (a PDE list). In the EM64T paging mode, the metadata includes, in one embodiment, vectors PTV, PDV, PDPV and PML4V, an active entry list (a PTE list), and active directory lists (an active PDE list, an active PDPE list and an active PML4E list).

One embodiment of a process for identifying recently modified entries of the guest translation data structure using metadata is discussed in more detail below in conjunction with FIG. 6.

At processing block **414**, processing logic synchronizes corresponding entries in the shadow translation data structure with the modified entries of the guest translation data structure. Accordingly, processing logic only needs to synchronize the entries that were modified, rather than re-populating the entire content of the shadow translation data structure.

In one embodiment extra storage is used to maintain some guest PD and/or PT contents as they were last synchronized. This permits the VMM to determine where modifications have been made without calculating or looking up additional relocation or permission information.

Note that certain modifications to the guest page tables do not require modifications to the shadow page tables. For example, if a guest PT contains a not present mapping which is subsequently modified, no change is required to the corresponding shadow PT.

FIGS. **5-14** illustrate various processes performed to support address translation in a virtual machine environment using the two-level hierarchy paging mode, according to different embodiments of the present invention. These processes may be performed by processing logic that may comprise hardware (e.g., dedicated logic, programmable logic, microcode, etc.), software (such as that run on a general purpose computer system or a dedicated machine), or a combination of both. In one embodiment, each of these processes is performed by an address translation module **126** of FIG. **1**.

FIG. **5** is a flow diagram of one embodiment of a process **500** for maintaining metadata for a shadow translation data structure such as a shadow PT hierarchy.

Referring to FIG. **5**, process **500** begins with processing logic creating a shadow page for each PD or PT page from the guest PT hierarchy (processing block **502**).

At processing block **504**, processing logic tracks page frames used as PDs or PTs in the guest PT hierarchy. In one embodiment, processing logic sets an entry in the PDV if a corresponding PFN is a PD in the guest PT hierarchy. Similarly, processing logic sets an entry in the PTV if a corresponding PFN is a PT in the guest PT hierarchy.

At processing block **506**, processing logic tracks mappings to any Dynamic Random Access Memory (DRAM) backed page (to identify pages that can potentially be PDs or PTs). In one embodiment, processing logic tracks mappings to DRAM based pages using an inverted page table (IPT) and an inverted page directory (IPD). The IPT is indexed by a PFN of a data page frame, with each entry containing a list of addresses of PTEs that map the data page frame. The IPD is indexed by a PFN of the page table, with each entry containing a list of addresses of PDEs that reference the PFN as a page table.

In one embodiment, at processing block **508**, processing logic identifies 4 MB pages in the guest PT hierarchy and creates a page table in the shadow PT hierarchy for each 4 MB page to avoid large page mappings and thereby reduce future synchronization time. Otherwise, an update of a 4 MB page would cause the synchronization of every PD and PT page within the 4 MB. In one embodiment, an inverted expansion table (IET) is used to track which PDEs in the guest PT hierarchy point to a 4 MB page. The IET is indexed by a PFN and attribute bits, with every entry listing PDEs that point to the exploded 4 MB page.

In an embodiment of the invention the IPD may be indexed by the address of the shadow PFN to minimize required address translation steps.

In IA32, memory type information (e.g., cacheability information) can be stored in PAT bits within the PDE/PTE that maps a page. This type information is not captured in a PDE that is a page-table pointer. Hence, if two 4 MB pages were to map the same region with different PAT attributes, then separate page tables would be required to convey the correct PAT attributes. Using separate expansion tables for each set of attributes resolves this issue.

At processing block **510**, processing logic identifies PTEs in the shadow PT hierarchy that map pages used as PD or PT in the guest PT hierarchy and creates an active PTE list.

At processing block **512**, processing logic identifies PDEs in the shadow PT hierarchy that point to PTs with PTEs identified in the active PTE list and creates an active PDE list.

13

Subsequently, at processing block **514**, if the guest OS modifies the structure of the guest PT hierarchy (e.g., adds or removes a PD or PT), processing logic changes the above active PTE and PDE lists accordingly.

FIG. **6** is a flow diagram of one embodiment of a process **600** for facilitating a change of an address space. Note that in IA32 the same CR3 value may also be reloaded to force a flush of stale TLB mappings. Similar processing steps are taken for a change of CR3 or for a CR3 reload.

Referring to FIG. **6**, process **600** begins with processing logic determining that a VM exit occurred due to a request of the VM to enable a different guest PT hierarchy (e.g., by issuing a CR3 load request) (processing block **602**).

In response, processing logic scans all active PDEs corresponding to the currently in-use shadow PT hierarchy identified in the active PDE list of the metadata to find which of these PDEs have been accessed (have an access attribute set to an access value) (processing block **604**), and then initializes the access attributes of the accessed PDEs (processing block **606**). In IA32, non-leaf paging tables do not support a dirty bit. If the accessed bit is clear, then no page within the 4 MB region has been read or written, so any guest page table or page directory cannot have been modified. However, the accessed bit does not distinguish between reads and writes, so 4 MB regions which have been accessed should be further processed even though it is possible that nothing has been modified. In architectures supporting a dirty bit for non-leaf page tables, the dirty bit is checked instead, and only regions which had been written to require further processing.

Next, for each accessed PDE, processing logic scans all shadow PTEs corresponding to the accessed active PDE in the active PTE list of the metadata to find which of these PTEs include mappings for an updated page (have an update attribute set to an update value) (processing block **608**).

Further, for each updated page, processing logic compares PD/PT entries in the guest PT hierarchy with corresponding entries in the shadow PT hierarchy (processing block **610**) and changes the corresponding entries of the shadow PT hierarchy to conform to the modified entries of the guest PT hierarchy (e.g., by removing from the shadow PT hierarchy a PTE/PDE absent in the guest PT hierarchy, by adding to the shadow PT hierarchy a new PTE/PDE recently added to the guest PT hierarchy, etc.) (processing block **612**). Note that adding PDEs may require the allocation and initialization of additional shadow PTs. This in turn may require updates to the various metadata structures maintained by the address translation module **126**.

At processing block **614**, processing logic initializes update attributes that were set to an update value. Updated mappings identify the pages that were modified by the guest OS.

At processing block **616**, processing logic synchronizes the shadow mappings based on modified guest pages and updates the metadata if needed due to the above modifications.

At processing logic **618**, processing logic determines whether a working set maintained by the VMM includes a shadow PD corresponding to the new guest PD requested by the VM. If so, processing logic requests the processor to load the base address of this shadow PT hierarchy (processing block **620**). If not, processing logic allocates a new shadow PT hierarchy corresponding to the requested guest PT hierarchy (processing block **622**), adds the PD of the new shadow PT hierarchy to the PDV (processing block **624**), adds each valid PDE to the PD of the new shadow PT hierarchy (processing block **626**), configures the active PDE and PTE lists to monitor the PTEs that map this PD for PD coverage (process-

14

ing block **628**), and then requests the processor to load the base address of this shadow PT hierarchy (processing block **620**). One embodiment of a process for monitoring a PTE is discussed in more detail below in conjunction with FIG. **13**.

FIG. **7** is a flow diagram of one embodiment of a process **700** for synchronizing entries of two translation data structures for a specified address. Process **700** may be performed, for example, as a result of the INVLPG instruction issued by the VM or as a result of an induced page fault.

Referring to FIG. **7**, process **700** begins with processing logic determining whether the mapping in the shadow PT hierarchy for the specified address is stale (i.e., there is valid mapping that does not correspond to the current contents of the guest page table) (processing box **702**). If not, processing logic proceeds to processing box **712**. If so, processing logic determines whether the stale entry mapped a PD or PT page (processing box **704**). If the stale entry did not map a PD or PT page, processing logic removes the stale entry (processing block **710**) and proceeds to processing box **712**.

If the stale entry did map a PD or PT page, processing logic further determines whether the mapped page has been updated (processing box **706**). If not, processing logic proceeds to processing block **710**. If so, processing logic updates, synchronizes, or removes the modified PD or PT shadow(s) (processing block **708**) and proceeds to processing block **710**. In one embodiment, the page is marked for future synchronization.

At processing box **712**, processing logic determines whether the guest PT hierarchy contains a new mapping for the specified address. If not, process **700** ends. If so, processing logic adds the new mapping as a corresponding PTE or PDE and, if necessary, creates a shadow page and updates the metadata according to the addition (processing block **714**).

FIG. **8** is a flow diagram of one embodiment of a process **800** for removing a shadow PT hierarchy from a working set of shadow PT hierarchies maintained by the VMM.

A shadow PT hierarchy may be removed from the working set upon detecting a deactivation of a corresponding process by the VM. The deactivation may be detected using heuristic defined for a relevant OS or employing a set of checks based on clues provided by the behavior of the guest VM with respect to the current address space. If the VM supports an interface through which the OS or a driver notifies the VMM of deactivations, then a heuristic may be avoided. A shadow PT hierarchy may also be removed due to resource constraints, e.g., because the amount of memory used for shadow structures exceeds a target threshold.

Referring to FIG. **8**, processing logic begins with removing each valid PDE from the PD in the shadow PT hierarchy (processing block **802**).

At processing block **804**, processing logic clears a corresponding entry in the PDV.

At processing block **806**, processing logic deallocates the PD page and removes the translation from a PD translation table (PDTT). The PDTT is used to track the address and type (e.g., PD or PT) of a page. The PDTT is indexed by a guest PFN, with each entry containing a physical PFN and metadata.

At processing block **808**, processing logic removes monitoring from the PTEs that map the PD. One embodiment of a process for removing monitoring from a PTE is discussed in more detail below in conjunction with FIG. **14**.

FIG. **9** is a flow diagram of one embodiment of a process **900** for adding an entry to a PD of a shadow PT hierarchy. For illustration, we will consider a present entry which maps a page table.

Referring to FIG. 9, processing logic begins with adding an entry for the PDE to the IPD (processing block 902).

At processing box 904, processing logic determines if the PT mapped by this PDE is set in the PTV. If so, the appropriate shadow PT is looked up in the PTTT (processing block 916), the new shadow PDE is created (processing block 914) and process 900 ends. If not, processing logic sets a corresponding vector in the PTV (processing block 906), allocates a shadow page and initializes the translation (processing block 908), populates the new shadow page table (processing block 910), updates active PTE/PDE lists and metadata to reflect that the guest page used as a page table by the current guest PDE is to be monitored (processing block 912), and adds the new PDE, adding it to the active PDE list if the shadow page table contains any active PTE list elements (processing block 914). One embodiment of a process for monitoring a PTE is discussed in more detail below in conjunction with FIG. 13.

FIG. 10 is a flow diagram of one embodiment of a process 1000 for removing an entry from a PD of a shadow PT hierarchy.

Referring to FIG. 10, processing logic begins with removing an entry for this PDE from the IPD PDE list (processing block 1002). If the PDE is in the active PDE list, then the active PDE list must be updated.

At processing box 1004, processing logic determines whether the PDE was the last entry to map the corresponding PT. If not, process 1000 ends. If so, processing logic clears the entry for the PT in the PTV (processing block 1006), removes each valid PTE (processing block 1008), updates the active PTE/PDE lists that map this PT for PT coverage (processing block 1010), and removes the shadow page translation and free the memory used to store the PT shadow page (processing block 1010).

FIG. 11 is a flow diagram of one embodiment of a process 1100 for adding an entry to a PT of a shadow PT hierarchy.

Referring to FIG. 11, processing logic begins with adding an entry for this PTE to the IPT (processing block 1102).

At processing box 1106, processing logic creates the shadow mapping and proceeds to processing box 1108.

At processing box 1108, processing logic determines whether a corresponding entry in the PDV or PTV is set. If not, process 1100 ends. If so, processing logic adds this entry to the active PTE list and updates associated metadata indicating if it maps a PD and/or PT page (processing block 1110). If the active PTE entry just created is the first for this page table, then the IPD must be consulted and each PDE which maps this page table page added to the active PDE list.

FIG. 12 is a flow diagram of one embodiment of a process 1200 for removing an entry from a PT of a shadow PT hierarchy.

Referring to FIG. 12, processing logic begins with determining whether this PTE maps a page set in the PDV or PTV (processing block 1202). If not, processing logic proceeds to processing block 1206. If so, processing logic removes the PTE from the active PTE list. If this was the last active PTE list element in the PT, then PDEs referencing this PT are removed from the active PDE list. (processing block 1204), and proceeds to processing block 1206.

At processing block 1206, processing logic removes the corresponding entry from the IPT.

FIG. 13 is a flow diagram of one embodiment of a process 1300 for monitoring a PTE of a shadow PT hierarchy. The steps shown in FIG. 13 represent the processing that may be required when the monitor recognizes that a page which has been mapped as a data page is being used as a page directory or page table page. This process will therefore be triggered by a status change for the page mapped by the PTE.

Referring to FIG. 13, processing logic begins with determining whether the PTE is identified in the active PTE list (processing box 1302). If so, processing logic adds the previously missing coverage (processing block 1304). This flow is triggered by a status change of the mapped page. Since this PTE was already in the active PTE list, it must be the case that this PTE was previously in use as a PT or PD, and is now in use in the other capacity as well. Such information may be explicitly stored with the entry or in associated metadata. If the PTE is not in the active PTE list, processing logic adds the PTE to the active PTE list and updates metadata accordingly (processing block 1306).

Next, at processing box 1308, processing logic determines whether the PTE is the first active PTE list entry for this PT. If not, process 1300 ends. If so, processing logic adds, to the active PDE list, entries that map this PT (as found through the IPD) (processing block 1310).

FIG. 14 is a flow diagram of one embodiment of a process 1400 for decreasing the monitoring coverage provided by a PTE of a shadow PT hierarchy. This process might be invoked when a process is removed from the working set, or the last PDE to reference a page table is removed, resulting in a change of status of a previously monitored page directory or page table page.

Referring to FIG. 14, processing logic begins with determining whether this PTE had monitored a page that was both a page table and page directory page (processing box 1402). If so, processing logic reduces the coverage level, indicating that the PTE now monitors a page as either a PT or as a PD, but not both (processing block 1404). If not, processing logic removes the PTE from the active PTE list (processing block 1406). Note that if the PTE was an element for a page tracked for its use in a single capacity, then it must now be the case that the page no longer requires monitoring.

Next, if the last active PTE list element in the PT was removed (processing box 1408), processing logic removes the corresponding entries which mapped this page table from the active PDE list (as found through the IPD) (processing block 1410).

As discussed above, the physical or virtual platform may comprise multiple processors. Each processor may in turn comprise one or more threads or logical processors. The processes discussed above can be used in a single-threaded system supporting a single-threaded VM or in a physical system with multiple logical processors that supports one or more VMs each containing a single virtual logical processor. Note that each VM has its own set of metadata, shadow page tables, etc. and that synchronization steps are confined to a given VM.

Thus, a method and apparatus for supporting address translation in a virtual machine environment have been described. It is to be understood that the above description is intended to be illustrative, and not restrictive. Many other embodiments will be apparent to those of skill in the art upon reading and understanding the above description. The scope of the invention should, therefore, be determined with reference to the appended claims, along with the full scope of equivalents to which such claims are entitled.

What is claimed is:

1. A method comprising:

receiving control transitioned from a virtual machine (VM) due to a privileged event pertaining to a translation-lookaside buffer (TLB);

determining which entries in a guest translation data structure were modified by the VM, based on metadata extracted from a shadow translation data structure main-

17

tained by a virtual machine monitor (VMM) and attributes associated with entries in the shadow translation data structure; and
synchronizing a first set of the entries in the shadow translation data structure with the modified entries in the guest translation data structure, without processing a second set of the entries from the shadow translation data structure,
wherein the first set of the entries in the shadow translation data structure includes entries that correspond to the modified entries from the guest translation data structure, and the second set of the entries in the shadow translation data structure includes entries that do not correspond to the modified entries from the guest translation data structure.

2. The method of claim 1 wherein:
the guest translation data structure is used by the VM for address translation operations; and
content of the shadow translation data structure is used by a processor to cache address translations in the TLB.

3. The method of claim 1 further comprising:
maintaining a working set of shadow translation data structures, each shadow translation data structure in the working set being associated with one of a plurality of active processes of the VM; and
reusing content of a shadow translation data structure from the working set that is associated with one of the plurality of active processes when a guest translation data structure associated with the one of the plurality of active processes is enabled.

4. The method of claim 3 wherein the privileged event is any one of a request of the VM to enable a different guest translation data structure, a page fault caused by one or more inconsistencies between entries of the guest translation data structure and entries of the shadow translation data structure, and a request of the VM to invalidate one or more address translations in a translation-lookaside buffer (TLB).

5. The method of claim 3 further comprising:
determining that the privileged event is any one of a page fault, an INVLPG instruction, and a request to enable a new guest translation data structure;
creating a new shadow translation data structure based on the new guest translation data structure; and
deriving metadata from the new shadow translation data structure.

6. The method of claim 3 further comprising:
determining that one of the plurality of the active processes of the VM is deactivated; and
removing a shadow translation data structure associated with the active process being deactivated from the working set.

7. The method of claim 3 further comprising:
determining one of the shadow translation data structures is no longer used; and
removing the one of the shadow translation data structures from the working set.

8. The method of claim 1 wherein the metadata includes a set of vectors, an active entry list, and one or more active directory lists.

9. The method of claim 8 wherein:
each vector in the set identifies frames used as pages at a corresponding level of the guest translation data structure;
the active entry list identifies mappings that map pages used in forming the guest translation data structure for which a shadow translation data structure exists; and

18

the one or more active directory lists identify higher level mapping structures referencing a lower level structure through which the shadow translation data structure can be accessed.

10. The method of claim 9 wherein:
the active entry list is an active page table (PT) entry list, the active PT entry (PTE) list identifying PTEs in the shadow translation data structure that map PT pages and PD pages from the guest translation data structure, and the one or more active directory lists include an active page directory (PD) entry list, the active PD entry (PDE) list identifying PDEs in the shadow translation data structure that point to PTs having the PTEs from the active PTE list.

11. The method of claim 9 wherein:
the active entry list is an active page table (PT) entry list, the active PT entry (PTE) list identifying mappings that map any of page map level 4 (PML4) pages, page directory pointer (PDP) pages, page directory (PD) pages, and PT pages; and
the one or more active directory lists include an active PD entry (PDE) list containing PDEs that reference a page with active PTE list entries, an active PDP entry (PDPE) list containing active PDPE entries which reference a PD with an active PDE list entry, and an active PML4E entry (PML4E) list containing entries which map a PDP with elements from the active PDPE list.

12. The method of claim 10 wherein the attributes associated with entries in the shadow translation data structure include access attributes associated with PDEs in the shadow translation data structure and update attributes associated with PTEs in the shadow translation data structure.

13. The method of claim 12 wherein determining which entries in the guest translation data structure were modified by the VM comprises:
identifying one or more PDEs from the active PDE list corresponding to the active PT hierarchy being synchronized that have access attributes set to an access value; and
for each of the identified PDEs, finding corresponding PTEs from the active PTE list that have update attributes set to an update value.

14. The method of claim 13 further comprising:
initializing the access attributes; and
initializing the update attributes.

15. An apparatus comprising:
a guest translation data structure used by a virtual machine (VM) for address translation operations;
a shadow translation data structure maintained by a virtual machine monitor (VMM), the shadow translation data structure deriving a format and structure from the guest translation data structure; and
an address translation module to determine, based on metadata extracted from the shadow translation data structure, which entries in the guest translation data structure were modified by the VM, and to synchronize a first set of the entries in the shadow translation data structure with the modified entries in the guest translation data structure, without processing a second set of the entries from the shadow translation data structure,
wherein the first set of the entries in the shadow translation data structure includes entries that correspond to the modified entries from the guest translation data structure, and the second set of the entries in the shadow translation data structure includes entries that do not correspond to the modified entries from the guest translation data structure.

19

16. The apparatus of claim 15 wherein the address translation module is further to maintain a working set of shadow translation data structures, each shadow translation data structure in the working set being associated with one of a plurality of active processes of the VM, and to reuse content of a shadow translation data structure from the working set that is associated with one of the plurality of active processes when a guest translation data structure associated with the one of the plurality of active processes is enabled.

17. The apparatus of claim 16 wherein the privileged event is any one of a request of the VM to enable a different guest translation data structure, a page fault caused by one or more inconsistencies between entries of the guest translation data structure and entries of the shadow translation data structure, and a request of the VM to invalidate one or more address translations in a translation-lookaside buffer (TLB).

18. A system comprising:

a dynamic random access memory (DRAM) to store a guest translation data structure used by a virtual machine (VM) for address translation operations, and a shadow translation data structure deriving a format and structure from the guest translation data structure; and

a processor, coupled to the DRAM, to determine, based on metadata extracted from the shadow translation data structure, which entries in the guest translation data structure were modified by the VM, and to synchronize a first set of the entries in the shadow translation data structure with the modified entries in the guest translation data structure, without processing a second set of the entries from the shadow translation data structure,

wherein the first set of the entries in the shadow translation data structure includes entries that correspond to the modified entries from the guest translation data structure, and the second set of the entries in the shadow translation data structure includes entries that do not correspond to the modified entries from the guest translation data structure.

19. The system of claim 18 wherein the processor is further to maintain a working set of shadow translation data structures, each shadow translation data structure in the working set being associated with one of a plurality of active processes of the VM, and to reuse content of a shadow translation data structure from the working set that is associated with one of the plurality of active processes when a guest translation data structure associated with the one of the plurality of active processes is enabled.

20. The system of claim 19 wherein the privileged event is any one of a request of the VM to enable a different guest translation data structure, a page fault caused by one or more

20

inconsistencies between entries of the guest translation data structure and entries of the shadow translation data structure, and a request of the VM to invalidate one or more address translations in a translation-lookaside buffer (TLB).

21. A computer-readable storage medium storing instructions which, when executed by a processing system, cause the processing system to perform a method, the method comprising:

receiving control transitioned from a virtual machine (VM) due to a privileged event pertaining to a translation-lookaside buffer (TLB);

determining which entries in a guest translation data structure were modified by the VM, based on metadata extracted from a shadow translation data structure maintained by a virtual machine monitor (VMM) and attributes associated with entries in the shadow translation data structure; and

synchronizing a first set of the entries in the shadow translation data structure with the modified entries in the guest translation data structure, without processing a second set of the entries from the shadow translation data structure,

wherein the first set of the entries in the shadow translation data structure includes entries that correspond to the modified entries from the guest translation data structure, and the second set of the entries in the shadow translation data structure includes entries that do not correspond to the modified entries from the guest translation data structure.

22. The computer-readable storage medium of claim 21 wherein the method further comprises:

maintaining a working set of shadow translation data structures, each shadow translation data structure in the working set being associated with one of a plurality of active processes of the VM, and

reusing content of a shadow translation data structure from the working set that is associated with one of the plurality of active processes when a guest translation data structure associated with the one of the plurality of active processes is enabled.

23. The computer-readable storage medium of claim 22 wherein the privileged event is any one of a request of the VM to enable a different guest translation data structure, a page fault caused by one or more inconsistencies between entries of the guest translation data structure and entries of the shadow translation data structure, and a request of the VM to invalidate one or more address translations in a translation-lookaside buffer (TLB).

* * * * *